



Red Hat Developer Toolset 6.1

User Guide

Installing and Using Red Hat Developer Toolset

Red Hat Developer Toolset 6.1 User Guide

Installing and Using Red Hat Developer Toolset

Jaromír Hradílek
Red Hat Customer Content Services

Matt Newsome
Red Hat Software Engineering

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

Vladimír Slávik
Red Hat Customer Content Services
vslavik@redhat.com

Legal Notice

Copyright © 2014-2017 Red Hat, Inc.

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

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

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

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

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

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

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

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

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

All other trademarks are the property of their respective owners.

Abstract

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 Red Hat Developer Toolset versions of the tools, and links to resources with more in-depth information.

Table of Contents

PART I. INTRODUCTION	5
CHAPTER 1. RED HAT DEVELOPER TOOLSET	6
1.1. ABOUT RED HAT DEVELOPER TOOLSET	6
1.2. MAIN FEATURES	8
1.3. COMPATIBILITY	8
1.4. GETTING ACCESS TO RED HAT DEVELOPER TOOLSET	9
1.5. INSTALLING RED HAT DEVELOPER TOOLSET	11
1.6. UPDATING RED HAT DEVELOPER TOOLSET	13
1.7. UNINSTALLING RED HAT DEVELOPER TOOLSET	14
1.8. USING RED HAT DEVELOPER TOOLSET CONTAINER IMAGES	14
1.9. ADDITIONAL RESOURCES	15
PART II. DEVELOPMENT TOOLS	17
CHAPTER 2. GNU COMPILER COLLECTION (GCC)	18
2.1. GNU C COMPILER	18
2.2. GNU C++ COMPILER	20
2.3. GNU FORTRAN COMPILER	23
2.4. ADDITIONAL RESOURCES	25
CHAPTER 3. GNU MAKE	27
3.1. INSTALLING MAKE	27
3.2. USING MAKE	27
3.3. USING MAKEFILES	28
3.4. ADDITIONAL RESOURCES	29
CHAPTER 4. BINUTILS	31
4.1. INSTALLING BINUTILS	32
4.2. USING THE GNU ASSEMBLER	32
4.3. USING THE GNU LINKER	32
4.4. USING OTHER BINARY TOOLS	33
4.5. ADDITIONAL RESOURCES	33
CHAPTER 5. ELFUTILS	35
5.1. INSTALLING ELFUTILS	35
5.2. USING ELFUTILS	36
5.3. ADDITIONAL RESOURCES	36
CHAPTER 6. DWZ	37
6.1. INSTALLING DWZ	37
6.2. USING DWZ	37
6.3. ADDITIONAL RESOURCES	37
PART III. DEBUGGING TOOLS	39
CHAPTER 7. GNU DEBUGGER (GDB)	40
7.1. INSTALLING THE GNU DEBUGGER	40
7.2. PREPARING A PROGRAM FOR DEBUGGING	40
7.3. RUNNING THE GNU DEBUGGER	41
7.4. LISTING SOURCE CODE	42
7.5. SETTING BREAKPOINTS	43
7.6. STARTING EXECUTION	45
7.7. DISPLAYING CURRENT VALUES	45

7.8. CONTINUING EXECUTION	45
7.9. ADDITIONAL RESOURCES	46
CHAPTER 8. STRACE	48
8.1. INSTALLING STRACE	48
8.2. USING STRACE	48
8.3. ADDITIONAL RESOURCES	51
CHAPTER 9. LTRACE	53
9.1. INSTALLING LTRACE	53
9.2. USING LTRACE	53
9.3. ADDITIONAL RESOURCES	56
CHAPTER 10. MEMSTOMP	57
10.1. INSTALLING MEMSTOMP	58
10.2. USING MEMSTOMP	58
10.3. ADDITIONAL RESOURCES	60
PART IV. PERFORMANCE MONITORING TOOLS	61
CHAPTER 11. SYSTEMTAP	62
11.1. INSTALLING SYSTEMTAP	62
11.2. USING SYSTEMTAP	63
11.3. ADDITIONAL RESOURCES	63
CHAPTER 12. VALGRIND	65
12.1. INSTALLING VALGRIND	65
12.2. USING VALGRIND	66
12.3. ADDITIONAL RESOURCES	66
CHAPTER 13. OPROFILE	68
13.1. INSTALLING OPROFILE	68
13.2. USING OPROFILE	69
13.3. ADDITIONAL RESOURCES	69
CHAPTER 14. DYNINST	71
14.1. INSTALLING DYNINST	71
14.2. USING DYNINST	71
14.3. ADDITIONAL RESOURCES	76
PART V. GETTING HELP	78
CHAPTER 15. ACCESSING RED HAT PRODUCT DOCUMENTATION	79
RED HAT DEVELOPER TOOLSET	79
RED HAT ENTERPRISE LINUX	79
CHAPTER 16. CONTACTING GLOBAL SUPPORT SERVICES	80
16.1. GATHERING REQUIRED INFORMATION	80
16.2. ESCALATING AN ISSUE	81
16.3. RE-OPENING A SERVICE REQUEST	82
16.4. ADDITIONAL RESOURCES	82
APPENDIX A. CHANGES IN VERSION 6.0	83
A.1. CHANGES IN BINUTILS	83
A.2. CHANGES IN ELFUTILS	86
A.3. CHANGES IN GCC	86
A.4. CHANGES IN MAKE	87

A.5. CHANGES IN GDB	87
A.6. CHANGES IN OPROFILE	89
A.7. CHANGES IN STRACE	90
A.8. CHANGES IN SYSTEMTAP	90
A.9. CHANGES IN DYNINST	91
A.10. CHANGES IN VALGRIND	91
APPENDIX B. CHANGES IN VERSION 6.1	93
B.1. CHANGES IN ELFUTILS	93
B.2. CHANGES IN GCC	93
B.3. CHANGES IN GDB	93
B.4. CHANGES IN DEVTOOLSET-6-TOOLCHAIN-RHEL7	93
B.5. CHANGES IN DEVTOOLSET-6-PERFTOOLS-RHEL7	94
APPENDIX C. REVISION HISTORY	95
INDEX	96

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. It 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, see [Section 1.3, “Compatibility”](#).

Red Hat Developer Toolset does not replace the default system tools provided with Red Hat Enterprise Linux 6 or 7 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.

What Is New in Red Hat Developer Toolset 6.1

Red Hat Developer Toolset 6.1 introduces support for additional 64-bit architectures:

- The 64-bit ARM architecture (AArch64)
- IBM POWER, big endian
- IBM POWER, little endian
- IBM z Systems



NOTE

The version number of Red Hat Developer Toolset has been raised from 4.1 to 6.0 with this release to signify the major improvements brought by **GCC 6.2.1** (upgraded from version 5.3.1 in the previous release of Red Hat Developer Toolset). There will be no Red Hat Developer Toolset 5.0.

Since Red Hat Developer Toolset 4.1, the Red Hat Developer Toolset content is also available in the ISO format at <https://access.redhat.com/downloads>, specifically for [Server](#) and [Workstation](#). Note that packages that require the *Optional* channel, which are discussed in [Section 1.5.3, “Installing Optional Packages”](#), cannot be installed from the ISO image.

Table 1.1. Red Hat Developer Toolset Components

Name	Version	Description
GCC	6.3.1	A portable compiler suite with support for C, C++, and Fortran.
binutils	2.27	A collection of binary tools and other utilities to inspect and manipulate object files and binaries.
elfutils	0.168	A collection of binary tools and other utilities to inspect and manipulate ELF files.

Name	Version	Description
dwz	0.12	A tool to optimize DWARF debugging information contained in ELF shared libraries and ELF executables for size.
GDB	7.12.1	A command line debugger for programs written in C, C++, and Fortran.
ltrace	0.7.91	A debugging tool to display calls to dynamic libraries that a program makes. It can also monitor system calls executed by programs.
strace	4.12	A debugging tool to monitor system calls that a program uses and signals it receives.
memstomp	0.1.5	A debugging tool to identify calls to library functions with overlapping memory regions that are not allowed by various standards.
SystemTap	3.0	A tracing and probing tool to monitor the activities of the entire system without the need to instrument, recompile, install, and reboot.
Valgrind	3.12.0	An instrumentation framework and a number of tools to profile applications in order to detect memory errors, identify memory management problems, and report any use of improper arguments in system calls.
OProfile	1.1.0	A system-wide profiler that uses the performance monitoring hardware on the processor to retrieve information about the kernel and executables on the system.
Dyninst	9.2.0	A library for instrumenting and working with user-space executables during their execution.
make	4.1	A dependency-tracking build automation tool.

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

1. 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”](#).
2. 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 are 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. A new major version of Red Hat Developer Toolset is 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, provides 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 2.2.4, “C++ 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 `sc1` utility.

1.2. MAIN FEATURES

The Red Hat Developer Toolset 6.1 brings the following changes:

- Red Hat Developer Toolset 6.1 adds support for Red Hat Enterprise Linux 6.9.
- The Red Hat Developer Toolset version of the **GNU Compiler Collection (GCC)** has been upgraded to version 6.3.1 with many bug fixes.
- The version of the **GNU Debugger (GDB)** included in Red Hat Developer Toolset provides several bug fixes.
- The `eu-readelf` utility from package `elfutils` now allows showing just the symbol table of the given ELF section name instead of symbols from all symbol tables.

For a full list of changes and features introduced in this release, see [Appendix B, *Changes in Version 6.1*](#).

1.3. COMPATIBILITY

Red Hat Developer Toolset 6.1 is available for Red Hat Enterprise Linux 6 and 7 for 64-bit Intel and AMD architectures. [Figure 1.1, “Red Hat Developer Toolset 6.1 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, see [Section 2.2.4, “C++ Compatibility”](#).

		Runs on				
		6.7	6.8	7.1	7.2	7.3
Built with	6.7	✓	✓	✓	✓	✓
	6.8	✗	✓	✓	✓	✓
	7.1	✗	✗	✓	✓	✓
	7.2	✗	✗	✗	✓	✓
	7.3	✗	✗	✗	✗	✓
		✗ Unsupported				✓ Supported

Figure 1.1. Red Hat Developer Toolset 6.1 Compatibility Matrix

1.4. GETTING ACCESS TO RED HAT DEVELOPER TOOLSET

Red Hat Developer Toolset is an offering that is distributed as a part of the Red Hat Software Collections content set, which is available to customers with Red Hat Enterprise Linux 6 and 7 subscriptions listed at <https://access.redhat.com/solutions/472793>. 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 Software Collections (and thus gain access to Red Hat Developer Toolset) using RHN Classic or Red Hat Subscription Management, see 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](#) collection of guides.

1.4.1. 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 Software Collections (which includes Red Hat Developer Toolset), and then enable that repository:

1. Determine the pool ID of a subscription that provides Red Hat Software Collections (and thus also Red Hat Developer Toolset). To do so, type the following at a shell prompt as **root** to display a list of all subscriptions that are available for your system:

```
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, see <https://access.redhat.com/solutions/472793>.

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

```
subscription-manager attach --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. Determine the exact name of the Red Hat Software Collections repository. To do so, type the following at a shell prompt as **root** to retrieve repository metadata and to display a list of available **Yum** repositories:

```
subscription-manager repos --list
```

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

```
rhel-variant-rhsc1-version-rpms  
rhel-variant-rhsc1-version-debug-rpms  
rhel-variant-rhsc1-version-source-rpms
```

In addition, certain packages, such as `devtoolset-6-gcc-plugin-devel`, depend on packages that are only available in the **Optional** channel. The repository names with these packages use the following format:

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

For both the regular repositories and optional repositories, replace *variant* with the Red Hat Enterprise Linux system variant (**server** or **workstation**), and *version* with the Red Hat Enterprise Linux system version (**6-eus**, **6**, or **7**).

4. Enable the repositories from step no. 3 by running the following command as **root**:

```
subscription-manager repos --enable repository
```

Replace *repository* with the name of the repository to enable.

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, see the [Red Hat Subscription Management](#) collection of guides.

1.4.2. Using RHN Classic

If you are running Red Hat Enterprise Linux 6, and your system is registered with RHN Classic, complete the following steps to subscribe to Red Hat Software Collections (which includes Red Hat Developer Toolset):

1. Determine the exact name of the Red Hat Software Collections channel. To do so, type the following at a shell prompt as **root** to display a list of all channels that are available to you:

```
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 following format:

```
rhel-x86_64-variant-version-rhsc1-2
```

In addition, certain packages, such as `devtoolset-6-gcc-plugin-devel`, depend on packages that are only available in the **Optional** channel. The name of this channel uses the following format:

```
rhel-x86_64-variant-optional-6
```

Replace *variant* with the Red Hat Enterprise Linux system variant (**server** or **workstation**).

2. Subscribe the system to the channels from step no. 1 by running the following command as **root**:

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

Replace *channel_name* with the name of the channel to enable.

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, see the [Red Hat Subscription Management](#) collection of guides.

1.5. INSTALLING RED HAT DEVELOPER TOOLSET

Red Hat Developer Toolset is distributed as a collection of RPM packages that can be installed, updated, uninstalled, and inspected by using the standard package management tools that are included in Red Hat Enterprise Linux. Note that a valid subscription that provides access to the Red Hat Software Collections content set is required in order to install Red Hat Developer Toolset on your system. For detailed instructions on how to associate your system with an appropriate subscription and get access to Red Hat Developer Toolset, see [Section 1.4, “Getting Access to Red Hat Developer Toolset”](#).



IMPORTANT

Before installing Red Hat Developer Toolset, install all available Red Hat Enterprise Linux updates.

1.5.1. Installing All Available Components

To install all components that are included in Red Hat Developer Toolset, install the `devtoolset-6` package by typing the following at a shell prompt as **root**:

`yum install devtoolset-6`

This installs all development, debugging, and performance monitoring tools, and other dependent packages to the system. Alternatively, you can choose to install only a selected package group as described in [Section 1.5.2, “Installing Individual Package Groups”](#).



NOTE

Note that since Red Hat Developer Toolset 3.0, the `scl-utils` package is not a part of Red Hat Developer Toolset, which is a change from preceding versions where the `scl` utility was installed along with the Red Hat Developer Toolset software collection.

1.5.2. Installing Individual Package Groups

To make it easier to install only certain components, such as the integrated development environment or the software development toolchain, Red Hat Developer Toolset is distributed with a number of meta packages that allow you to install selected package groups as described in [Table 1.2, “Red Hat Developer Toolset Meta Packages”](#).

Table 1.2. Red Hat Developer Toolset Meta Packages

Package Name	Description	Installed Components
<code>devtoolset-6-perftools</code>	Performance monitoring tools	SystemTap, Valgrind, OProfile, Dyninst
<code>devtoolset-6-toolchain</code>	Development and debugging tools	GCC, make, GDB, binutils, elfutils, dwz, memstomp, strace, ltrace

To install any of these meta packages, type the following at a shell prompt as **root**:

```
yum install package_name
```

Replace *package_name* with a space-separated list of meta packages you want to install. For example, to install only the development and debugging toolchain and packages that depend on it, type as **root**:

```
~]# yum install devtoolset-6-toolchain
```

Alternatively, you can choose to install all available components as described in [Section 1.5.1, “Installing All Available Components”](#).

1.5.3. Installing Optional Packages

Red Hat Developer Toolset is distributed with a number of optional packages that are not installed by default. To list all Red Hat Developer Toolset packages that are available to you but not installed on your system, type the following command at a shell prompt:

```
yum list available devtoolset-6-\*
```

To install any of these optional packages, run as **root**:


```
yum install package_name
```

Replace *package_name* with a space-separated list of packages that you want to install. For example, to install the `devtoolset-6-gdb-gdbserver` and `devtoolset-6-gdb-doc` packages, type:

```
~]# yum install devtoolset-6-gdb-gdbserver devtoolset-6-gdb-doc
```

1.5.4. Installing Debugging Information

To install debugging information for any of the Red Hat Developer Toolset packages, make sure that the `yum-utils` package is installed and run the following command as **root**:

```
debuginfo-install package_name
```

For example, to install debugging information for the `devtoolset-6-dwz` package, type:

```
~]# debuginfo-install devtoolset-6-dwz
```

Note that in order to use this command, you need to have access to the repository with these packages. If your system is registered with Red Hat Subscription Management, enable the **`rhel-variant-rhsc1-version-debug-rpms`** repository as described in [Section 1.4.1, “Using Red Hat Subscription Management”](#). If your system is registered with RHN Classic, subscribe the system to the **`rhel-x86_64-variant-version-debuginfo`** channel as described in [Section 1.4.2, “Using RHN Classic”](#). For more information on how to get access to debuginfo packages, see <https://access.redhat.com/site/solutions/9907>.



NOTE

The `devtoolset-6-package_name-debuginfo` packages can conflict with the corresponding packages from the base Red Hat Enterprise Linux system or from other versions of Red Hat Developer Toolset. This conflict also occurs in a multilib environment, where 64-bit debuginfo packages conflict with 32-bit debuginfo packages. Manually uninstall the conflicting debuginfo packages prior to installing Red Hat Developer Toolset 6.1 and install only relevant debuginfo packages when necessary.

1.6. UPDATING RED HAT DEVELOPER TOOLSET

1.6.1. Updating to a Minor Version

When a new minor version of Red Hat Developer Toolset is available, 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 development, debugging, and performance monitoring tools, and other dependent packages.



IMPORTANT

Use of Red Hat Developer Toolset requires the removal of any earlier pre-release versions of it. Additionally, it is not possible to update to Red Hat Developer Toolset 6.1 from a pre-release version of Red Hat Developer Toolset, including beta releases. If you have previously installed any pre-release version of Red Hat Developer Toolset, uninstall it from your system as described in [Section 1.7, “Uninstalling Red Hat Developer Toolset”](#) and install the new version as documented in [Section 1.5, “Installing Red Hat Developer Toolset”](#).

1.6.2. Updating to a Major Version

When a new major version of Red Hat Developer Toolset is available, you can install it in parallel with the previous version. For detailed instructions on how to install Red Hat Developer Toolset on your system, see [Section 1.5, “Installing Red Hat Developer Toolset”](#).

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-6\* libasan libatomic libcilkrts libitm liblsan  
libtsan libubsan
```

This removes the **GNU Compiler Collection**, **GNU Debugger**, **binutils**, and other packages that are a part of Red Hat Developer Toolset from the system.



NOTE

Red Hat Developer Toolset 6.1 for Red Hat Enterprise Linux 7 no longer includes the **libatomic** and **libitm** libraries, which the above command attempts to remove, because they are not required for a proper function of Red Hat Developer Toolset components on that system. Nevertheless, the above command works as expected even on Red Hat Enterprise Linux 7.

Note that the 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. USING RED HAT DEVELOPER TOOLSET CONTAINER IMAGES

Docker-formatted container images can be used to run Red Hat Developer Toolset components inside virtual software containers, thus isolating them from the host system and allowing for their rapid deployment. For detailed description of the Red Hat Developer Toolset docker-formatted container images and Red Hat Developer Toolset *Dockerfiles*, see [Using Red Hat Software Collections Container Images](#).



NOTE

The **Docker** package, which contains the **Docker** daemon, command-line tool, and other necessary components for building and using docker-formatted container images, is currently available only for the Server variant of the Red Hat Enterprise Linux 7 product. Docker-formatted container images cannot be run on Red Hat Enterprise Linux 6, and Red Hat Developer Toolset Dockerfiles are not distributed for Red Hat Enterprise Linux 6.

Follow the instructions outlined at [Getting Docker in RHEL 7](#) to set up an environment for building and using docker-formatted container images.

1.9. ADDITIONAL RESOURCES

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

Online Documentation

- [Red Hat Subscription Management](#) collection of guides — The Red Hat Subscription Management collection of guides provides detailed information on how to manage subscriptions on Red Hat Enterprise Linux.
- [Red Hat Developer Toolset 6.1 Release Notes](#) — The *Release Notes* for Red Hat Developer Toolset 6.1 contain more information.
- [Red Hat Enterprise Linux 6 Developer Guide](#) and [Red Hat Enterprise Linux 7 Developer Guide](#) — The *Developer Guides* for Red Hat Enterprise Linux 6 and 7 provide more information on the **Eclipse** IDE, libraries and runtime support, compiling and building, debugging, and profiling on these systems.
- [Red Hat Enterprise Linux 6 Installation Guide](#) and [Red Hat Enterprise Linux 7 Installation Guide](#) — The *Installation Guides* for Red Hat Enterprise Linux 6 and 7 explain 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 7 System Administrator's Guide](#) — The *System Administrator's Guide* for Red Hat Enterprise Linux 7 documents relevant information regarding the deployment, configuration, and administration of Red Hat Enterprise Linux 7.
- [Using Red Hat Software Collections Container Images](#) — This book provides information on how to use container images based on Red Hat Software Collections. The available container images include applications, daemons, databases, as well as the Red Hat Developer Toolset container images. The images can be run on Red Hat Enterprise Linux 7 Server and Red Hat Enterprise Linux Atomic Host.
- [Get Started with Docker Formatted Container Images on Red Hat Systems](#) — The guide contains a comprehensive overview of information about building and using docker-formatted container images on Red Hat Enterprise Linux 7 and Red Hat Enterprise Linux Atomic.

See Also

- [Appendix B, Changes in Version 6.1](#) provides a list of changes and improvements over the version of the GNU Compiler Collection and GNU Debugger in the previous version of Red Hat Developer Toolset.

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 6.3.1**. This version is more recent than the version included in Red Hat Enterprise Linux and provides a number of bug fixes and enhancements.

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-6-gcc` package and is automatically installed with `devtoolset-6-toolchain` 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:

```
sc1 enable devtoolset-6 '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:


```
sc1 enable devtoolset-6 '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:

```
sc1 enable devtoolset-6 'gcc -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 `gcc` as default:

```
sc1 enable devtoolset-6 '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.

**NOTE**

The Red Hat Developer Toolset 6.1 version of **GCC** supports Cilk+, an extension to the C and C++ languages for parallel programming, which can be enabled using the **-fcilkplus** option. A runtime library, **libcilkrts**, is included in this release to support Cilk+. The **libcilkrts** library has been a part of Red Hat Enterprise Linux since version 7.2, but the package is not included in all supported Red Hat Enterprise Linux releases. To enable dynamic linkage of binaries and libraries built with Red Hat Developer Toolset 6.1 **GCC** using Cilk+ features on supported Red Hat Enterprise Linux releases that do not contain **libcilkrts**, install the **libcilkrts.so** shared library from Red Hat Developer Toolset 6.1 with such binaries or libraries.

Example 2.1. Compiling a C Program on the Command Line

Consider 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-6 '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-6-gcc-c++` package and is automatically installed with the `devtoolset-6-toolchain` 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-6 '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-6 '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-6 '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:


```
scl enable devtoolset-6 '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

Consider 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-6 '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!
```

2.2.4. C++ Compatibility

All compilers from Red Hat Enterprise Linux versions 5, 6, and 7 and from Red Hat Developer Toolset versions 1, 2, 3, 4, and 6 in any `-std` mode are compatible with any other of those compilers in C++98 mode.

A compiler in C++11 or C++14 mode is only guaranteed to be compatible with another compiler in C++11 or C++14 mode if they are from the same release series (for example from Red Hat Developer Toolset 6.x).



IMPORTANT

All compatibility information mentioned in this section is relevant only for Red Hat-supplied versions of the GCC C++ compiler.

2.2.4.1. C++ ABI

Any C++98-compliant binaries or libraries built explicitly with `-std=c++98` or `-std=gnu++98` can be freely mixed with binaries and shared libraries built by the Red Hat Enterprise Linux 5, 6 or 7 system toolchain GCC. Red Hat recommends use of the `-std=c++98` or `-std=gnu++98` modes for production software development.

The default language standard setting for Red Hat Developer Toolset is C++14 with GNU extensions, equivalent to explicitly using option `-std=gnu++14`.

Using and mixing the C++11 and C++14 language versions is supported in Red Hat Developer Toolset only when all C++ objects compiled with the respective flag have been built using the same major version of GCC. 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 options `-std=c++11`, `-std=gnu++11`, `-std=c++14`, and `-std=gnu++14` options cannot be guaranteed, and so is not supported.



IMPORTANT

Use of C++11 and C++14 features in your application requires careful consideration of the above ABI compatibility information.

The mixing of objects, binaries and libraries, built by the Red Hat Enterprise Linux 6 or 7 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` or `-std=c++14` or `-std=gnu++14` flags using the GCC in Red Hat Developer Toolset is explicitly not supported.

Aside from the C++11 and C++14 ABI, discussed above, [the Red Hat Enterprise Linux Application Compatibility Specification](#) is unchanged for Red Hat Developer Toolset. When mixing objects built with Red Hat Developer Toolset with those built with the Red Hat Enterprise Linux 6 or 7 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 Red Hat 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.

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-6-gcc-gfortran` package and is automatically installed with `devtoolset-6-toolchain` 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-6 '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-6 '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-6 '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-6 '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

Consider 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-6 '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, see 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-6 '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-6 '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-6-libstdc++-docs
```

Once installed, HTML documentation is available at `/opt/rh/devtoolset-6/root/usr/share/doc/devtoolset-6-libstdc++-docs-6.3.1/html/index.html`.

Online Documentation

- [Red Hat Enterprise Linux 6 Developer Guide](#) and [Red Hat Enterprise Linux 7 Developer Guide](#) — The *Developer Guides* for Red Hat Enterprise Linux 6 and 7 provide 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

- [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 5, *elfutils*](#) explains how to use **elfutils**, a collection of binary tools to inspect and manipulate ELF files.
- [Chapter 6, *dwz*](#) explains how to use **dwz** to optimize DWARF debugging information contained in ELF shared libraries and ELF executables for size.
- [Chapter 7, *GNU Debugger \(GDB\)*](#) provides information on how to debug programs written in C, C++, and Fortran.

CHAPTER 3. GNU MAKE

The **GNU make** utility, commonly abbreviated **make**, is a tool for controlling the generation of executables from source files. **make** automatically determines which parts of a complex program have changed and need to be recompiled. **make** uses configuration files called *Makefiles* to control the way programs are built.

Red Hat Developer Toolset is distributed with **make 4.1**. This version is more recent than the version included in Red Hat Enterprise Linux and provides a number of bug fixes and enhancements.

3.1. INSTALLING MAKE

In Red Hat Developer Toolset, **GNU make** is provided by the `devtoolset-6-make` package and is automatically installed with `devtoolset-6-toolchain` as described in [Section 1.5, “Installing Red Hat Developer Toolset”](#).

3.2. USING MAKE

To build a program without using a Makefile, run the **make** tool as follows:

```
sc1 enable devtoolset-6 'make source_file_without_extension'
```

This command makes use of implicit rules that are defined for a number of programming languages, including C, C++, and Fortran. The result is a binary file named *source_file_without_extension* in the current working directory.

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 **make** as default:

```
sc1 enable devtoolset-6 'bash'
```

NOTE

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

```
which make
```

Red Hat Developer Toolset's **make** 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 **make**:

```
make -v
```

Example 3.1. Building a C Program Using make

Consider 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 build this source code using the implicit rules defined by the **make** utility from Red Hat Developer Toolset, type:

```
~]$ scl enable devtoolset-6 'make hello'
cc      hello.c  -o hello
```

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

3.3. USING MAKEFILES

To build complex programs that consist of a number of source files, **make** uses configuration files called *Makefiles* that control how to compile the components of a program and build the final executable. Makefiles can also contain instructions for cleaning the working directory, installing and uninstalling program files, and other operations.

make automatically uses files named **GNUmakefile**, **makefile**, or **Makefile** in the current directory. To specify another file name, use:

```
make -f make_file
```

Describing the details of Makefile syntax is beyond the scope of this guide. See [GNU make](#), the official **GNU make** manual, which provides an in-depth description of the **GNU make** utility, Makefile syntax, and their usage.

The full **make** manual is also available in the Texinfo format as a part of your installation. To view this manual, type:

```
~]$ scl enable devtoolset-6 'info make'
```

Example 3.2. Building a C Program Using a Makefile

Consider the following universal Makefile named **Makefile** for building the simple C program introduced in [Example 3.1, “Building a C Program Using make”](#). The Makefile defines some variables and specifies four *rules*, which consist of *targets* and their *recipes*. Note that the lines with recipes must start with the TAB character:

```
CC=gcc
CFLAGS=-c -Wall
SOURCE=hello.c
OBJ=$(SOURCE:.c=.o)
EXE=hello

all: $(SOURCE) $(EXE)

$(EXE): $(OBJ)
    $(CC) $(OBJ) -o $@

.c.o:
```



```

$(CC) $(CFLAGS) $< -o $@

clean:
    rm -rf $(OBJ) $(EXE)

```

To build the **hello.c** program using this Makefile, run the **make** utility:

```

~]$ scl enable devtoolset-6 'make'
gcc -c -Wall hello.c -o hello.o
gcc hello.o -o hello

```

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

To clean the working directory, run:

```

~]$ scl enable devtoolset-6 'make clean'
rm -rf hello.o hello

```

This removes the object and binary files from the working directory.

3.4. ADDITIONAL RESOURCES

A detailed description of the **GNU make** tool and its features is beyond the scope of this book. For more information, see the resources listed below.

Installed Documentation

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

```
scl enable devtoolset-6 'man make'
```

- The full **make** manual, which includes detailed information about Makefile syntax, is also available in the Texinfo format. To display the info manual for the version included in Red Hat Developer Toolset, type:

```
scl enable devtoolset-6 'info make'
```

Online Documentation

- [GNU make](#) — The official **GNU make** manual provides an in-depth description of the **GNU make** utility, Makefile syntax, and their usage.

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 \(GCC\)](#) explains how to use the **GNU Compiler Collection**, a portable compiler suite with support for a wide selection of programming languages.

- [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 6, *dwz*](#) explains how to use **dwz** to optimize DWARF debugging information contained in ELF shared libraries and ELF executables for size.
- [Chapter 7, *GNU Debugger \(GDB\)*](#) provides information on how to debug programs written in C, C++, and Fortran.

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. See [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 2.27**. This version is more recent than the version included in Red Hat Enterprise Linux and the previous release of Red Hat Developer Toolset and provides bug fixes and enhancements.

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

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

Name	Description
strip	Discards all symbols from object files.

4.1. INSTALLING BINUTILS

In Red Hat Developer Toolset, **binutils** are provided by the `devtoolset-6-binutils` package and are automatically installed with `devtoolset-6-toolchain` 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:

```
sc1 enable devtoolset-6 '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 **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 **as** as default:

```
sc1 enable devtoolset-6 '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:

```
sc1 enable devtoolset-6 '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 **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 **ld** as default:

```
sc1 enable devtoolset-6 '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 an assembler. For a complete list of these tools, see [Table 4.1, “Tools Included in binutils for Red Hat Developer Toolset”](#).

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

```
sc1 enable devtoolset-6 'tool [option...] file_name'
```

See [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:

```
sc1 enable devtoolset-6 '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:

```
sc1 enable devtoolset-6 '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, see the resources listed below.

Installed Documentation

- `as(1)`, `ld(1)`, `addr2line(1)`, `ar(1)`, `c++filt(1)`, `dwp(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:

```
❯ scli enable devtoolset-6 'man tool'
```

Online Documentation

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

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 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. See [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 0.168**. This version is more recent than the version included the previous release of Red Hat Developer Toolset and provides some bug fixes and enhancements.

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

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

5.1. INSTALLING ELFUTILS

In Red Hat Developer Toolset, **elfutils** is provided by the `devtoolset-6-elfutils` package and is automatically installed with `devtoolset-6-toolchain` 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-6 'tool [option...] file_name'
```

See [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-6 '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:

```
sc1 enable devtoolset-6 '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, see 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 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 6, *dwz*](#) explains how to use **dwz** to optimize DWARF debugging information contained in ELF shared libraries and ELF executables for size.

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 0.12**.

6.1. INSTALLING DWZ

In Red Hat Developer Toolset, the **dwz** utility is provided by the `devtoolset-6-dwz` package and is automatically installed with `devtoolset-6-toolchain` 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-6 '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-6 '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, see 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:

```
sc1 enable devtoolset-6 '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. DEBUGGING TOOLS

CHAPTER 7. 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 7.12.1**. This version is more recent than the version included in Red Hat Enterprise Linux and the previous release of Red Hat Developer Toolset and provides some enhancements and numerous bug fixes.

7.1. INSTALLING THE GNU DEBUGGER

In Red Hat Developer Toolset, the **GNU Debugger** is provided by the `devtoolset-6-gdb` package and is automatically installed with `devtoolset-6-toolchain` as described in [Section 1.5, “Installing Red Hat Developer Toolset”](#).

7.2. 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:

```
sc1 enable devtoolset-6 'gcc -g -o output_file input_file...'
```

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

```
sc1 enable devtoolset-6 'g++ -g -o output_file input_file...'
```

Example 7.1. Compiling a C Program With Debugging Information

Consider a source file named `fibonacci.c` that has 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-6 '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 7.2. 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 0:2.17-105.el7 will be installed
...
```

7.3. RUNNING THE GNU DEBUGGER

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

```
scl enable devtoolset-6 '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 **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 **gdb** as default:

```
scl enable devtoolset-6 '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 7.3. Running the **gdb** Utility on the **fibonacci** Binary File

Assuming that you have successfully compiled the **fibonacci** binary file as shown in [Example 7.1, “Compiling a C Program With Debugging Information”](#), you can start debugging it with **gdb** by typing the following at a shell prompt:

```
~]$ scl enable devtoolset-6 'gdb fibonacci'
GNU gdb (GDB) Red Hat Enterprise Linux (7.12.1-47.el7)
Copyright (C) 2015 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/>.
Find the GDB manual and other documentation resources online at:
<http://www.gnu.org/software/gdb/documentation/>.
For help, type "help".
Type "apropos word" to search for commands related to "word".
(gdb)
```

7.4. 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 the 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 7.4. Listing the Source Code of the fibonacci Binary File

The `fibonacci.c` file listed in [Example 7.1, “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     }
```

7.5. SETTING BREAKPOINTS

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 7.5. Setting a New Breakpoint

Assuming that you have compiled the `fibonacci.c` file listed in [Example 7.1, “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 7.6. Listing Breakpoints

Assuming that you have followed the instructions in [Example 7.5, “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   0x000000000004004e5 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
```

Example 7.7. Deleting an Existing Breakpoint

Assuming that you have compiled the `fibonacci.c` file listed in [Example 7.1, “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
```


7.6. STARTING EXECUTION

To start an 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 the first breakpoint (if any) is reached, when an error occurs, or when the program terminates.

Example 7.8. Executing the fibonacci Binary File

Assuming that you have followed the instructions in [Example 7.5, “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=0x7fffffff4d8) at fibonacci.c:10
10         printf("%ld ", b);
```

7.7. 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 7.9. Displaying the Current Values of Variables

Assuming that you have followed the instructions in [Example 7.8, “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
```

7.8. 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 7.10. Continuing the Execution of the `fibonacci` Binary File

Assuming that you have followed the instructions in [Example 7.8, “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 the 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
```

7.9. ADDITIONAL RESOURCES

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

Online Documentation

- [Red Hat Enterprise Linux 6 Developer Guide](#) and [Red Hat Enterprise Linux 7 Developer Guide](#) — The *Developer Guides* for Red Hat Enterprise Linux 6 and 7 provide 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

- [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 8, *strace*](#) documents how to use the **strace** utility to monitor system calls that a program uses and signals it receives.
- [Chapter 10, *memstomp*](#) documents how to use the **memstomp** utility to identify calls to library functions with overlapping memory regions that are not allowed by various standards.

CHAPTER 8. STRACE

strace is a diagnostic and debugging tool for the command line that can be used to trace system calls that are made and received by a running process. It records the name of each system call, its arguments, and its return value, as well as signals received by the process and other interactions with the kernel, and prints this record to standard error output or a selected file.

Red Hat Developer Toolset is distributed with **strace 4.12**.

8.1. INSTALLING STRACE

In Red Hat Enterprise Linux, the **strace** utility is provided by the `devtoolset-6-strace` package and is automatically installed with `devtoolset-6-toolchain` as described in [Section 1.5, “Installing Red Hat Developer Toolset”](#).

8.2. USING STRACE

To run the **strace** utility on a program you want to analyze, type the following at a shell prompt:

```
sc1 enable devtoolset-6 'strace program [argument...]
```

Replace *program* with the name of the program you want to analyze, and *argument* with any command line options and arguments you want to supply to this program. Alternatively, you can run the utility on an already running process by using the `-p` command line option followed by the process ID:

```
sc1 enable devtoolset-6 'strace -p process_id'
```

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 **strace** as default:

```
sc1 enable devtoolset-6 'bash'
```

NOTE

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

```
which strace
```

Red Hat Developer Toolset's **strace** 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 **strace**:

```
strace -V
```

8.2.1. Redirecting Output to a File

By default, **strace** prints the name of each system call, its arguments and the return value to standard error output. To redirect this output to a file, use the `-o` command line option followed by the file name:

-

```
scl enable devtoolset-6 'strace -o file_name program [argument...]
```

Replace *file_name* with the name of the file.

Example 8.1. Redirecting Output to a File

Consider a slightly modified version of the `fibonacci` file from [Example 7.1, “Compiling a C Program With Debugging Information”](#). This executable file displays the Fibonacci sequence and optionally allows you to specify how many members of this sequence to list. To run the `strace` utility on this file and redirect the trace output to `fibonacci.log`, type:

```
~]$ scl enable devtoolset-6 'strace -o fibonacci.log ./fibonacci 20'
1 1 2 3 5 8 13 21 34 55 89 144 233 377 610 987 1597 2584 4181 6765
```

This creates a new plain-text file called `fibonacci.log` in the current working directory.

8.2.2. Tracing Selected System Calls

To trace only a selected set of system calls, run the `strace` utility with the `-e` command line option:

```
scl enable devtoolset-6 'strace -e expression program [argument...]
```

Replace *expression* with a comma-separated list of system calls to trace or any of the keywords listed in [Table 8.1, “Commonly Used Values of the -e Option”](#). For a detailed description of all available values, see the `strace(1)` manual page.

Table 8.1. Commonly Used Values of the -e Option

Value	Description
<code>file</code>	System calls that accept a file name as an argument.
<code>process</code>	System calls that are related to process management.
<code>network</code>	System calls that are related to networking.
<code>signal</code>	System calls that are related to signal management.
<code>ipc</code>	System calls that are related to inter-process communication (IPC).
<code>desc</code>	System calls that are related to file descriptors.

Example 8.2. Tracing Selected System Calls

Consider the `employee` file from [Example 10.1, “Using memstomp”](#). To run the `strace` utility on this executable file and trace only the `mmap` and `munmap` system calls, type:

```
~]$ scl enable devtoolset-6 'strace -e mmap,munmap ./employee'
mmap(NULL, 4096, PROT_READ|PROT_WRITE, MAP_PRIVATE|MAP_ANONYMOUS, -1, 0)
```

```

= 0x7f896c744000
mmap(NULL, 61239, PROT_READ, MAP_PRIVATE, 3, 0) = 0x7f896c735000
mmap(0x3146a00000, 3745960, PROT_READ|PROT_EXEC,
MAP_PRIVATE|MAP_DENYWRITE, 3, 0) = 0x3146a00000
mmap(0x3146d89000, 20480, PROT_READ|PROT_WRITE,
MAP_PRIVATE|MAP_FIXED|MAP_DENYWRITE, 3, 0x189000) = 0x3146d89000
mmap(0x3146d8e000, 18600, PROT_READ|PROT_WRITE,
MAP_PRIVATE|MAP_FIXED|MAP_ANONYMOUS, -1, 0) = 0x3146d8e000
mmap(NULL, 4096, PROT_READ|PROT_WRITE, MAP_PRIVATE|MAP_ANONYMOUS, -1, 0)
= 0x7f896c734000
mmap(NULL, 4096, PROT_READ|PROT_WRITE, MAP_PRIVATE|MAP_ANONYMOUS, -1, 0)
= 0x7f896c733000
mmap(NULL, 4096, PROT_READ|PROT_WRITE, MAP_PRIVATE|MAP_ANONYMOUS, -1, 0)
= 0x7f896c732000
munmap(0x7f896c735000, 61239) = 0
mmap(NULL, 4096, PROT_READ|PROT_WRITE, MAP_PRIVATE|MAP_ANONYMOUS, -1, 0)
= 0x7f896c743000
John,john@example.comDoe,
+++ exited with 0 +++

```

8.2.3. Displaying Time Stamps

To prefix each line of the trace with the exact time of the day in hours, minutes, and seconds, run the **strace** utility with the **-t** command line option:

```
scl enable devtoolset-6 'strace -t program [argument...]
```

To also display milliseconds, supply the **-t** option twice:

```
scl enable devtoolset-6 'strace -tt program [argument...]
```

To prefix each line of the trace with the time required to execute the respective system call, use the **-r** command line option:

```
scl enable devtoolset-6 'strace -r program [argument...]
```

Example 8.3. Displaying Time Stamps

Consider an executable file named **pwd**. To run the **strace** utility on this file and include time stamps in the output, type:

```

~]$ scl enable devtoolset-6 'strace -tt pwd'
19:43:28.011815 execve("./pwd", ["/.pwd"], [/* 36 vars */]) = 0
19:43:28.012128 brk(0) = 0xcd3000
19:43:28.012174 mmap(NULL, 4096, PROT_READ|PROT_WRITE,
MAP_PRIVATE|MAP_ANONYMOUS, -1, 0) = 0x7fc869cb0000
19:43:28.012427 open("/etc/ld.so.cache", O_RDONLY) = 3
19:43:28.012446 fstat(3, {st_mode=S_IFREG|0644, st_size=61239, ...}) = 0
19:43:28.012464 mmap(NULL, 61239, PROT_READ, MAP_PRIVATE, 3, 0) =
0x7fc869ca1000

```

```
19:43:28.012483 close(3)          = 0
...
19:43:28.013410 +++ exited with 0 +++
```

8.2.4. Displaying a Summary

To display a summary of how much time was required to execute each system call, how many times were these system calls executed, and how many errors were encountered during their execution, run the **strace** utility with the **-c** command line option:

```
scl enable devtoolset-6 'strace -c program [argument...]
```

Example 8.4. Displaying a Summary

Consider an executable file named **lsblk**. To run the **strace** utility on this file and display a trace summary, type:

```
~]$ scl enable devtoolset-6 'strace -c lsblk > /dev/null'
% time      seconds  usecs/call   calls   errors syscall
-----
 80.88      0.000055      1         106      16  open
 19.12      0.000013      0         140           munmap
  0.00      0.000000      0         148           read
  0.00      0.000000      0           1           write
  0.00      0.000000      0         258           close
  0.00      0.000000      0          37           2 stat
...
-----
100.00      0.000068                   1790          35 total
```

8.3. ADDITIONAL RESOURCES

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

Installed Documentation

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

```
scl enable devtoolset-6 'man strace'
```

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 9, ltrace](#) provides information on how to trace program library calls using the **ltrace** tool.
- [Chapter 7, GNU Debugger \(GDB\)](#) provides information on how to debug programs written in C, C++, and Fortran.

- [Chapter 10, *memstomp*](#) documents how to use the **memstomp** utility to identify calls to library functions with overlapping memory regions that are not allowed by various standards.

CHAPTER 9. LTRACE

ltrace is a diagnostic and debugging tool for the command line that can be used to display calls that are made to shared libraries. It uses the dynamic library hooking mechanism, which prevents it from tracing calls to statically linked libraries. **ltrace** also displays return values of the library calls. The output is printed to standard error output or to a selected file.

Red Hat Developer Toolset is distributed with **ltrace 0.7.91**. While the base version **ltrace** remains the same as in the previous release of Red Hat Developer Toolset, various enhancements and bug fixes have ported.

9.1. INSTALLING LTRACE

In Red Hat Enterprise Linux, the **ltrace** utility is provided by the `devtoolset-6-ltrace` package and is automatically installed with `devtoolset-6-toolchain` as described in [Section 1.5, “Installing Red Hat Developer Toolset”](#).

9.2. USING LTRACE

To run the **ltrace** utility on a program you want to analyze, type the following at a shell prompt:

```
sc1 enable devtoolset-6 'ltrace program [argument...]
```

Replace *program* with the name of the program you want to analyze, and *argument* with any command line options and arguments you want to supply to this program. Alternatively, you can run the utility on an already running process by using the `-p` command line option followed by the process ID:

```
sc1 enable devtoolset-6 'ltrace -p process_id'
```

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 **ltrace** as default:

```
sc1 enable devtoolset-6 'bash'
```

NOTE

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

```
which ltrace
```

Red Hat Developer Toolset's **ltrace** 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 **ltrace**:

```
ltrace -v
```

9.2.1. Redirecting Output to a File

By default, **ltrace** prints the name of each system call, its arguments and the return value to standard error output. To redirect this output to a file, use the **-o** command line option followed by the file name:

```
scl enable devtoolset-6 'ltrace -o file_name program [argument...]
```

Replace *file_name* with the name of the file.

Example 9.1. Redirecting Output to a File

Consider a slightly modified version of the **fibonacci** file from [Example 7.1, “Compiling a C Program With Debugging Information”](#). This executable file displays the Fibonacci sequence and optionally allows you to specify how many members of this sequence to list. To run the **ltrace** utility on this file and redirect the trace output to **fibonacci.log**, type:

```
~]$ scl enable devtoolset-6 'ltrace -o fibonacci.log ./fibonacci 20'
1 1 2 3 5 8 13 21 34 55 89 144 233 377 610 987 1597 2584 4181 6765
```

This creates a new plain-text file called **fibonacci.log** in the current working directory.

9.2.2. Tracing Selected Library Calls

To trace only a selected set of library calls, run the **ltrace** utility with the **-e** command line option:

```
scl enable devtoolset-6 'ltrace -e expression program [argument...]
```

Replace *expression* with a chain of rules to specify the library calls to trace. The rules can consist of patterns that identify symbol names (such as **malloc** or **free**) and patterns that identify library SONAMEs (such as **libc.so**). For example, to trace call to the **malloc** and **free** function but to omit those that are done by the **libc** library, use:

```
scl enable devtoolset-6 'ltrace -e malloc+free-@libc.so* program'
```

Example 9.2. Tracing Selected Library Calls

Consider the **ls** command. To run the **ltrace** utility on this program and trace only the **opendir**, **readdir**, and **closedir** function calls, type:

```
~]$ scl enable devtoolset-6 'ltrace -e opendir+readdir+closedir ls'
ls->opendir(".") = { 3 }
ls->readdir({ 3 }) = { 61533, "." }
ls->readdir({ 3 }) = { 131, ".." }
ls->readdir({ 3 }) = { 67185100, "BUILDROOT" }
ls->readdir({ 3 }) = { 202390772, "SOURCES" }
ls->readdir({ 3 }) = { 60249, "SPECS" }
ls->readdir({ 3 }) = { 67130110, "BUILD" }
ls->readdir({ 3 }) = { 136599168, "RPMS" }
ls->readdir({ 3 }) = { 202383274, "SRPMS" }
ls->readdir({ 3 }) = nil
ls->closedir({ 3 }) = 0
BUILD BUILDROOT RPMS SOURCES SPECS SRPMS
+++ exited (status 0) +++
```

For a detailed description of available filter expressions, see the `ltrace(1)` manual page.

9.2.3. Displaying Time Stamps

To prefix each line of the trace with the exact time of the day in hours, minutes, and seconds, run the `ltrace` utility with the `-t` command line option:

```
scl enable devtoolset-6 'ltrace -t program [argument...]'
```

To also display milliseconds, supply the `-t` option twice:

```
scl enable devtoolset-6 'ltrace -tt program [argument...]'
```

To prefix each line of the trace with the time required to execute the respective system call, use the `-r` command line option:

```
scl enable devtoolset-6 'ltrace -r program [argument...]'
```

Example 9.3. Displaying Time Stamps

Consider the `pwd` command. To run the `ltrace` utility on this program and include time stamps in the output, type:

```
~]$ scl enable devtoolset-6 'ltrace -tt pwd'
13:27:19.631371 __libc_start_main([ "pwd" ] <unfinished ...>
13:27:19.632240 getenv("POSIXLY_CORRECT") = nil
13:27:19.632520 strrchr("pwd", '/') = nil
13:27:19.632786 setlocale(LC_ALL, "") =
"en_US.UTF-8"
13:27:19.633220 bindtextdomain("coreutils", "/usr/share/locale") =
"/usr/share/locale"
13:27:19.633471 textdomain("coreutils") =
"coreutils"
...
13:27:19.637110 +++ exited (status 0) +++
```

9.2.4. Displaying a Summary

To display a summary of how much time was required to execute each system call and how many times were these system calls executed, run the `ltrace` utility with the `-c` command line option:

```
scl enable devtoolset-6 'ltrace -c program [argument...]'
```

Example 9.4. Displaying a Summary

Consider the `lsblk` command. To run the `ltrace` utility on this program and display a trace summary, type:

```
~]$ scl enable devtoolset-6 'ltrace -c lsblk > /dev/null'
```

% time	seconds	usecs/call	calls	function
53.60	0.261644	261644	1	__libc_start_main
4.48	0.021848	58	374	mbrtowc
4.41	0.021524	57	374	wcwidth
4.39	0.021409	57	374	__ctype_get_mb_cur_max
4.38	0.021359	57	374	iswprint
4.06	0.019838	74	266	readdir64
3.21	0.015652	69	224	strlen
...				
100.00	0.488135		3482	total

9.3. ADDITIONAL RESOURCES

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

Installed Documentation

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

```
scl enable devtoolset-6 'man ltrace'
```

Online Documentation

- [ltrace for RHEL 6 and 7](#) — This article on the Red Hat Developer Blog offers additional in-depth information (including practical examples) on how to use **ltrace** for application debugging.

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, strace](#) provides information on how to trace program system calls using the **strace** tool.
- [Chapter 7, GNU Debugger \(GDB\)](#) provides information on how to debug programs written in C, C++, and Fortran.
- [Chapter 10, memstomp](#) documents how to use the **memstomp** utility to identify calls to library functions with overlapping memory regions that are not allowed by various standards.

CHAPTER 10. MEMSTOMP

memstomp is a command line tool that can be used to identify function calls with overlapping memory regions in situations when such an overlap is not permitted by various standards. It intercepts calls to the library functions listed in [Table 10.1, “Function Calls Inspected by memstomp”](#) and for each memory overlap, it displays a detailed backtrace to help you debug the problem.

Similarly to **Valgrind**, the **memstomp** utility inspects applications without the need to recompile them. However, it is much faster than this tool and therefore serves as a convenient alternative to it.

Red Hat Developer Toolset is distributed with **memstomp 0.1.5**.

Table 10.1. Function Calls Inspected by memstomp

Function	Description
memcpy	Copies n bytes from one memory area to another and returns a pointer to the second memory area.
memccpy	Copies a maximum of n bytes from one memory area to another and stops when a certain character is found. It either returns a pointer to the byte following the last written byte, or NULL if the given character is not found.
mempcpy	Copies n bytes from one memory area to another and returns a pointer to the byte following the last written byte.
strcpy	Copies a string from one memory area to another and returns a pointer to the second string.
stpcpy	Copies a string from one memory area to another and returns a pointer to the terminating null byte of the second string.
strncpy	Copies a maximum of n characters from one string to another and returns a pointer to the second string.
stpncpy	Copies a maximum of n characters from one string to another. It either returns a pointer to the terminating null byte of the second string, or if the string is not null-terminated, a pointer to the byte following the last written byte.
strcat	Appends one string to another while overwriting the terminating null byte of the second string and adding a new one at its end. It returns a pointer to the new string.
strncat	Appends a maximum of n characters from one string to another while overwriting the terminating null byte of the second string and adding a new one at its end. It returns a pointer to the new string.
wmemcpy	The wide-character equivalent of the memcpy() function that copies n wide characters from one array to another and returns a pointer to the second array.

Function	Description
wmemcpy	The wide-character equivalent of the memcpy() function that copies <i>n</i> wide characters from one array to another and returns a pointer to the byte following the last written wide character.
wscpy	The wide-character equivalent of the strcpy() function that copies a wide-character string from one array to another and returns a pointer to the second array.
wcsncpy	The wide-character equivalent of the strncpy() function that copies a maximum of <i>n</i> wide characters from one array to another and returns a pointer to the second string.
wscat	The wide-character equivalent of the strcat() function that appends one wide-character string to another while overwriting the terminating null byte of the second string and adding a new one at its end. It returns a pointer to the new string.
wcsncat	The wide-character equivalent of the strncat() function that appends a maximum of <i>n</i> wide characters from one array to another while overwriting the terminating null byte of the second wide-character string and adding a new one at its end. It returns a pointer to the new string.

10.1. INSTALLING MEMSTOMP

In Red Hat Developer Toolset, the **memstomp** utility is provided by the `devtoolset-6-memstomp` package and is automatically installed with `devtoolset-6-toolchain` as described in [Section 1.5, “Installing Red Hat Developer Toolset”](#).

10.2. USING MEMSTOMP

To run the **memstomp** utility on a program you want to analyze, type the following at a shell prompt:

```
scl enable devtoolset-6 'memstomp program [argument...]
```

To immediately terminate the analyzed program when a problem is detected, run the utility with the **--kill** (or **-k** for short) command line option:

```
scl enable devtoolset-6 'memstomp --kill program [argument...]
```

The use of the **--kill** option is especially recommended if you are analyzing a multi-threaded program; the internal implementation of backtraces is not thread-safe and running the **memstomp** utility on a multi-threaded program without this command line option can therefore produce unreliable results.

Additionally, if you have compiled the analyzed program with the debugging information or this debugging information is available to you, you can use the **--debug-info** (or **-d**) command line option to produce a more detailed backtrace:

```
scl enable devtoolset-6 'memstomp --debug-info program [argument...]
```

For detailed instructions on how to compile your program with the debugging information built in the binary file, see [Section 7.2, “Preparing a Program for Debugging”](#). For information on how to install debugging information for any of the Red Hat Developer Toolset packages, see [Section 1.5.4, “Installing Debugging Information”](#).

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 `memstomp` as default:

```
scl enable devtoolset-6 'bash'
```

Example 10.1. Using memstomp

In the current working directory, create a source file named `employee.c` with the following contents:

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

#define BUFSIZE 80

int main(int argc, char *argv[]) {
    char employee[BUFSIZE] = "John,Doe,john@example.com";
    char name[BUFSIZE] = {0};
    char surname[BUFSIZE] = {0};
    char *email;
    size_t length;

    /* Extract the information: */
    memccpy(name, employee, ',', BUFSIZE);
    length = strlen(name);
    memccpy(surname, employee + length, ',', BUFSIZE);
    length += strlen(surname);
    email = employee + length;

    /* Compose the new entry: */
    strcat(employee, surname);
    strcpy(employee, name);
    strcat(employee, email);

    /* Print the result: */
    puts(employee);

    return 0;
}
```

Compile this program into a binary file named `employee` by using the following command:

```
~]$ scl enable devtoolset-6 'gcc -rdynamic -g -o employee employee.c'
```

To identify erroneous function calls with overlapping memory regions, type:

```
~]$ scl enable devtoolset-6 'memstomp --debug-info ./employee'
memstomp: 0.1.4 successfully initialized for process employee (pid
14887).
```

```
strcat(dest=0x7fff13afc265, src=0x7fff13afc269, bytes=21) overlap for
employee(14887)
  ??:0   strcpy()
  ??:0   strcpy()
  ??:0   _Exit()
  ??:0   strcat()
employee.c:26  main()
  ??:0   __libc_start_main()
  ??:0   _start()
John, john@example.comDoe,
```

10.3. ADDITIONAL RESOURCES

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

Installed Documentation

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

```
scl enable devtoolset-6 'man memstomp'
```

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, *GNU Debugger \(GDB\)*](#) provides information on how to debug programs written in C, C++, and Fortran.
- [Chapter 8, *strace*](#) documents how to use the `strace` utility to monitor system calls that a program uses and signals it receives.
- [Chapter 12, *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. PERFORMANCE MONITORING TOOLS

CHAPTER 11. 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 3.0**. This version is more recent than the version included in the previous release of Red Hat Developer Toolset and provides numerous bug fixes and enhancements.

Table 11.1. Tools Distributed with SystemTap for Red Hat Developer Toolset

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

11.1. INSTALLING SYSTEMTAP

In Red Hat Developer Toolset, **SystemTap** is provided by the `devtoolset-6-systemtap` package and is automatically installed with `devtoolset-6-perftools` as described in [Section 1.5, “Installing Red Hat Developer Toolset”](#).

**NOTE**

The Red Hat Developer Toolset version of **SystemTap** is available for both Red Hat Enterprise Linux 6 and Red Hat Enterprise Linux 7, but some new features are only offered by the Red Hat Developer Toolset version of **SystemTap** for Red Hat Enterprise Linux 7.

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:

```
sc1 enable devtoolset-6 '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, see the [Red Hat Enterprise Linux 6 SystemTap Beginners Guide](#) or the [Red Hat Enterprise Linux 7 SystemTap Beginners Guide](#).

11.2. USING SYSTEMTAP

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

```
sc1 enable devtoolset-6 'tool [option...]
```

See [Table 11.1, “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-6 '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-6 '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
```

11.3. ADDITIONAL RESOURCES

A detailed description of **SystemTap** and its features is beyond the scope of this book. For more information, see 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:

```
scl enable devtoolset-6 '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:

```
scl enable devtoolset-6 'man staprun'
```

- *SystemTap Tapset Reference Manual* — HTML documentation on the most common tapset definitions is located at `/opt/rh/devtoolset-6/root/usr/share/doc/devtoolset-6-systemtap-client-2.8/index.html`.

Online Documentation

- [Red Hat Enterprise Linux 6 SystemTap Beginners Guide](#) and [Red Hat Enterprise Linux 7 SystemTap Beginners Guide](#) — The *SystemTap Beginners Guides* for Red Hat Enterprise Linux 6 and 7 provide an introduction to **SystemTap** and its usage.
- [Red Hat Enterprise Linux 6 SystemTap Tapset Reference](#) and [Red Hat Enterprise Linux 7 SystemTap Tapset Reference](#) — The *SystemTap Tapset Reference* for Red Hat Enterprise Linux 6 and 7 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 12, 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 13, OProfile](#) explains how to use **OProfile** to determine which sections of code consume the greatest amount of CPU time and why.
- [Chapter 14, Dyninst](#) documents how to use the Dyninst library to instrument a user-space executable.

CHAPTER 12. VALGRIND

Valgrind is an instrumentation framework that ships with a number of tools for profiling applications. It can be used to detect various memory errors and memory-management problems, such as the use of uninitialized memory or an 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**, see [Table 12.1, “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 3.12.0**. This version is more recent than the version included in the previous release of Red Hat Developer Toolset and provides numerous bug fixes and enhancements.

Table 12.1. Tools Distributed with Valgrind for Red Hat Developer Toolset

Name	Description
Memcheck	Detects memory management problems by intercepting system calls and checking all read and write operations.
Cachegrind	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).
Callgrind	Generates a call graph representing the function call history.
Helgrind	Detects synchronization errors in multithreaded C, C++, and Fortran programs that use POSIX threading primitives.
DRD	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.
Massif	Monitors heap and stack usage.

12.1. INSTALLING VALGRIND

In Red Hat Developer Toolset, **Valgrind** is provided by the `devtoolset-6-valgrind` package and is automatically installed with `devtoolset-6-perftools`.

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



NOTE

Note that if you use **Valgrind** in combination with the **GNU Debugger**, it is recommended that you use the version of **GDB** that is included in Red Hat Developer Toolset to ensure that all features are fully supported.

12.2. USING VALGRIND

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

```
sc1 enable devtoolset-6 'valgrind [--tool=tool] program [argument...]
```

See [Table 12.1, “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:

```
sc1 enable devtoolset-6 'valgrind --tool=cachegrind program [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 **Valgrind** as default:

```
sc1 enable devtoolset-6 '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
```

12.3. ADDITIONAL RESOURCES

A detailed description of **Valgrind** and its features is beyond the scope of this book. For more information, see 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:

```
sc1 enable devtoolset-6 'man valgrind'
```

- *Valgrind Documentation* — HTML documentation for **Valgrind** is located at `/opt/rh/devtoolset-6/root/usr/share/doc/devtoolset-6-valgrind-3.9.0/html/index.html`.

Online Documentation

- [Red Hat Enterprise Linux 6 Developer Guide](#) and [Red Hat Enterprise Linux 7 Developer Guide](#) — The *Developer Guides* for Red Hat Enterprise Linux 6 and 7 provide more information about **Valgrind** and its **Eclipse** plug-in.
- [Red Hat Enterprise Linux 6 Performance Tuning Guide](#) [Red Hat Enterprise Linux 7 Performance Tuning Guide](#) — The *Performance Tuning Guides* for Red Hat Enterprise Linux 6 and 7 provide 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 10, *memstomp*](#) documents how to use the **memstomp** utility to identify calls to library functions with overlapping memory regions that are not allowed by various standards.
- [Chapter 11, *SystemTap*](#) provides an introduction to **SystemTap** and explains how to use it to monitor the activities of a running system.
- [Chapter 13, *OProfile*](#) explains how to use **OProfile** to determine which sections of code consume the greatest amount of CPU time and why.
- [Chapter 14, *Dyninst*](#) documents how to use the Dyninst library to instrument a user-space executable.

CHAPTER 13. 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**, see [Table 13.1, “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 it 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 1.1.0**.

Table 13.1. Tools Distributed with OProfile for Red Hat Developer Toolset

Name	Description
oprofiled	The OProfile daemon that collects profiling data.
operf	Intended to replace the deprecated opcontrol tool. The operf tool uses the Linux Performance Events subsystem, allowing users to target their profiling more precisely, as a single process or system-wide, and allowing OProfile to co-exist better with other tools using the performance monitoring hardware on your system. Unlike opcontrol , no initial setup is required, and it can be used without the root privileges unless the --system-wide option is in use.
opannotate	Generates an annotated source file or assembly listing from the profiling data.
oparchive	Generates a directory containing executable, debug, and sample files.
opgprof	Generates a summary of a profiling session in a format compatible with gprof .
ophelp	Displays a list of available events.
opimport	Converts a sample database file from a foreign binary format to the native format.
opjitconv	Converts a just-in-time (JIT) dump file to the Executable and Linkable Format (ELF).
opreport	Generates image and symbol summaries of a profiling session.
ocount	A new tool for counting the number of times particular events occur during the duration of a monitored command.

13.1. INSTALLING OPROFILE

In Red Hat Developer Toolset, **OProfile** is provided by the `devtoolset-6-oprofile` package and is automatically installed with `devtoolset-6-perftools` as described in [Section 1.5, “Installing Red Hat Developer Toolset”](#).

13.2. 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-6 'tool [option...]
```

See [Table 13.1, “Tools Distributed with OProfile for Red Hat Developer Toolset”](#) for a list of tools that are distributed with **OProfile**. For example, to use the `ophelp` command to list available events in the XML format, type:

```
sc1 enable devtoolset-6 '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-6 'bash'
```

NOTE

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

```
which operf
```

Red Hat Developer Toolset's `operf` 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**:

```
operf --version
```

13.3. ADDITIONAL RESOURCES

A detailed description of **OProfile** and its features is beyond the scope of this book. For more information, see 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-6 'man oprofile'
```

- `opannotate(1)`, `oparchive(1)`, `operf(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-6 'man tool'
```

Online Documentation

- [Red Hat Enterprise Linux 6 Developer Guide](#) and [Red Hat Enterprise Linux 7 Developer Guide](#) — The *Developer Guides* for Red Hat Enterprise Linux 6 and 7 provide 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 7 System Administrator's Guide](#) — The *System Administrator's Guide* for Red Hat Enterprise Linux 7 documents how to use the **operf** tool.

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 11, *SystemTap*](#) provides an introduction to **SystemTap** and explains how to use it to monitor the activities of a running system.
- [Chapter 12, *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 14, *Dyninst*](#) documents how to use the Dyninst library to instrument a user-space executable.

CHAPTER 14. DYNINST

The **Dyninst** library provides an *application programming interface* (API) for instrumenting and working with user-space executables during their execution. It can be used to insert code into a running program, change certain subroutine calls, or even remove them from the program. It serves as a valuable debugging and performance-monitoring tool. The **Dyninst** API is also commonly used along with **SystemTap** to allow non-**root** users to instrument user-space executables.

Red Hat Developer Toolset is distributed with **Dyninst 9.2.0**.

14.1. INSTALLING DYNINST

In Red Hat Developer Toolset, the **Dyninst** library is provided by the `devtoolset-6-dyninst` package and is automatically installed with `devtoolset-6-perftools` as described in [Section 1.5, “Installing Red Hat Developer Toolset”](#). In addition, it is recommended that you also install the **GNU Compiler Collection** provided by the `devtoolset-6-toolchain` package.

If you intend to write a custom instrumentation for binaries, install the relevant header files by running the following command as **root**:

```
yum install devtoolset-6-dyninst-devel
```

You can also install API documentation for this library by typing the following at a shell prompt as **root**:

```
yum install devtoolset-6-dyninst-doc
```

For a complete list of documents that are included in the `devtoolset-6-dyninst-doc` package, see [Section 14.3, “Additional Resources”](#). For detailed instructions on how to install optional packages to your system, see [Section 1.5, “Installing Red Hat Developer Toolset”](#).

14.2. USING DYNINST

14.2.1. Using Dyninst with SystemTap

To use **Dyninst** along with **SystemTap** to allow non-**root** users to instrument user-space executables, run the `stap` command with the `--dyninst` (or `--runtime=dyninst`) command line option. This tells `stap` to translate a **SystemTap** script into C code that uses the **Dyninst** library, compile this C code into a shared library, and then load the shared library and run the script. Note that when executed like this, the `stap` command also requires the `-c` or `-x` command line option to be specified.

To use the **Dyninst** runtime to instrument an executable file, type the following at a shell prompt:

```
sc1 enable devtoolset-6 "stap --dyninst -c 'command' [option...]  
[argument...]"
```

Similarly, to use the **Dyninst** runtime to instrument a user's process, type:

```
sc1 enable devtoolset-6 "stap --dyninst -x process_id [option...]  
[argument...]"
```

See [Chapter 11, *SystemTap*](#) for more information about the Red Hat Developer Toolset version of **SystemTap**. For a general introduction to **SystemTap** and its usage, see the [SystemTap Beginners Guide](#) for Red Hat Enterprise Linux 6 or the [SystemTap Beginners Guide](#) for Red Hat Enterprise Linux 7.

Example 14.1. Using Dyninst with SystemTap

Consider a source file named `exercise.C` that has the following contents:

```
#include <stdio.h>

void print_iteration(int value) {
    printf("Iteration number %d\n", value);
}

int main(int argc, char **argv) {
    int i;
    printf("Enter the starting number: ");
    scanf("%d", &i);
    for(; i>0; --i)
        print_iteration(i);
    return 0;
}
```

This program prompts the user to enter a starting number and then counts down to 1, calling the `print_iteration()` function for each iteration in order to print the number to the standard output. To compile this program on the command line using the `g++` compiler from Red Hat Developer Toolset, type the following at a shell prompt:

```
~]$ scl enable devtoolset-6 'g++ -g -o exercise exercise.C'
```

Now consider another source file named `count.stp` with the following contents:

```
#!/usr/bin/stap

global count = 0

probe process.function("print_iteration") {
    count++
}

probe end {
    printf("Function executed %d times.\n", count)
}
```

This **SystemTap** script prints the total number of times the `print_iteration()` function was called during the execution of a process. To run this script on the `exercise` binary file, type:

```
~]$ scl enable devtoolset-6 "stap --dyninst -c './exercise' count.stp"
Enter the starting number: 5
Iteration number 5
Iteration number 4
Iteration number 3
```

```
Iteration number 2
Iteration number 1
Function executed 5 times.
```

14.2.2. Using Dyninst as a Stand-alone Application

Before using the **Dyninst** library as a stand-alone application, set the value of the **DYNINSTAPI_RT_LIB** environment variable to the path to the runtime library file. You can do so by typing the following at a shell prompt:

```
export DYNINSTAPI_RT_LIB=/opt/rh/devtoolset-
6/root/usr/lib64/dyninst/libdyninstAPI_RT.so
```

This sets the **DYNINSTAPI_RT_LIB** environment variable in the current shell session.

[Example 14.2, “Using Dyninst as a Stand-alone Application”](#) illustrates how to write and build a program to monitor the execution of a user-space process. For a detailed explanation of how to use **Dyninst**, see the resources listed in [Section 14.3, “Additional Resources”](#).

Example 14.2. Using Dyninst as a Stand-alone Application

Consider the **exercise.C** source file from [Example 14.1, “Using Dyninst with SystemTap”](#): this program prompts the user to enter a starting number and then counts down to 1, calling the **print_iteration()** function for each iteration in order to print the number to standard output.

Now consider another source file named **count.C** with the following contents:

```
#include <stdio.h>
#include <fcntl.h>
#include "BPatch.h"
#include "BPatch_process.h"
#include "BPatch_function.h"
#include "BPatch_Vector.h"
#include "BPatch_thread.h"
#include "BPatch_point.h"

void usage() {
    fprintf(stderr, "Usage: count <process_id> <function>\n");
}

// Global information for counter
BPatch_variableExpr *counter = NULL;

void createCounter(BPatch_process *app, BPatch_image *appImage) {
    int zero = 0;
    counter = app->malloc(*appImage->findType("int"));
    counter->writeValue(&zero);
}

bool interceptfunc(BPatch_process *app,
                  BPatch_image *appImage,
                  char *funcName) {
    BPatch_Vector<BPatch_function *> func;
```

```

appImage->findFunction(funcName, func);
if(func.size() == 0) {
    fprintf(stderr, "Unable to find function to instrument()\n");
    exit (-1);
}
BPatch_Vector<BPatch_snippet *> incCount;
BPatch_Vector<BPatch_point *> *points;
points = func[0]->findPoint(BPatch_entry);
if ((*points).size() == 0) {
    exit (-1);
}

BPatch_arithExpr counterPlusOne(BPatch_plus, *counter,
BPatch_constExpr(1));
BPatch_arithExpr addCounter(BPatch_assign, *counter, counterPlusOne);

return app->insertSnippet(addCounter, *points);
}

void printCount(BPatch_thread *thread, BPatch_exitType) {
    int val = 0;
    counter->readValue(&val, sizeof(int));
    fprintf(stderr, "Function executed %d times.\n", val);
}

int main(int argc, char *argv[]) {
    int pid;
    BPatch bpatch;
    if (argc != 3) {
        usage();
        exit(1);
    }
    pid = atoi(argv[1]);
    BPatch_process *app = bpatch.processAttach(NULL, pid);
    if (!app) exit (-1);
    BPatch_image *appImage = app->getImage();
    createCounter(app, appImage);
    fprintf(stderr, "Finding function %s(): ", argv[2]);
    BPatch_Vector<BPatch_function*> countFuncs;
    fprintf(stderr, "OK\nInstrumenting function %s(): ", argv[2]);
    interceptfunc(app, appImage, argv[2]);
    bpatch.registerExitCallback(printCount);
    fprintf(stderr, "OK\nWaiting for process %d to exit...\n", pid);
    app->continueExecution();
    while (!app->isTerminated())
        bpatch.waitForStatusChange();
    return 0;
}

```

Note that a client application is expected to destroy all **BPatch** objects before any of the **Dyninst** library destructors are called. Otherwise the mutator might terminate unexpectedly with a segmentation fault. To work around this problem, set the **BPatch** object of the mutator as a local variable in the **main()** function. Or, if you need to use **BPatch** as a global variable, manually detach all the mutatee processes before the mutator exits.

This program accepts a process ID and a function name as command line arguments and then prints the total number of times the function was called during the execution of the process. You can use the following **Makefile** to build these two files:

```
DTS      = /opt/rh/devtoolset-4/root
CXXFLAGS = -g -I$(DTS)/usr/include/dyninst
LBITS    := $(shell getconf LONG_BIT)

ifeq ($(LBITS),64)
    DYNINSTLIBS = $(DTS)/usr/lib64/dyninst
else
    DYNINSTLIBS = $(DTS)/usr/lib/dyninst
endif

.PHONY: all
all: count exercise

count: count.C
g++ $(CXXFLAGS) count.C -I /usr/include/dyninst -c
g++ $(CXXFLAGS) count.o -L $(DYNINSTLIBS) -ldyninstAPI -o count

exercise: exercise.C
g++ $(CXXFLAGS) exercise.C -o exercise

.PHONY: clean
clean:
rm -rf *~ *.o count exercise
```

To compile the two programs on the command line using the **g++** compiler from Red Hat Developer Toolset, run the **make** utility as follows:

```
~]$ scl enable devtoolset-6 make
g++ -g -I/opt/rh/devtoolset-6/root/usr/include/dyninst count.C -c
g++ -g -I/opt/rh/devtoolset-6/root/usr/include/dyninst count.o -L
/opt/rh/devtoolset-6/root/usr/lib64/dyninst -ldyninstAPI -o count
g++ -g -I/opt/rh/devtoolset-6/root/usr/include/dyninst exercise.C -o
exercise
```

This creates new binary files called **exercise** and **count** in the current working directory.

In one shell session, execute the **exercise** binary file as follows and wait for it to prompt you to enter the starting number:

```
~]$ ./exercise
Enter the starting number:
```

Do not enter this number. Instead, start another shell session and type the following at its prompt to set the **DYNINSTAPI_RT_LIB** environment variable and execute the **count** binary file:

```
~]$ export DYNINSTAPI_RT_LIB=/opt/rh/devtoolset-
6/root/usr/lib64/dyninst/libdyninstAPI_RT.so
~]$ ./count `pidof exercise` print_iteration
```

```
Finding function print_iteration(): OK
Instrumenting function print_iteration(): OK
Waiting for process 8607 to exit...
```

Now switch back to the first shell session and enter the starting number as requested by the **exercise** program. For example:

```
Enter the starting number: 5
Iteration number 5
Iteration number 4
Iteration number 3
Iteration number 2
Iteration number 1
```

When the **exercise** program terminates, the **count** program displays the number of times the **print_iteration()** function was executed:

```
Function executed 5 times.
```

14.3. ADDITIONAL RESOURCES

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

Installed Documentation

The devtoolset-6-dyninst-doc package installs the following documents in the `/opt/rh/devtoolset-6/root/usr/share/doc/devtoolset-6-dyninst-doc-8.2.1/` directory:

- *Dyninst Programmer's Guide* — A detailed description of the Dyninst API is stored in the **DyninstAPI.pdf** file.
- *DynC API Programmer's Guide* — An introduction to DynC API is stored in the **dync_API.pdf** file.
- *ParseAPI Programmer's Guide* — An introduction to the ParseAPI is stored in the **ParseAPI.pdf** file.
- *PatchAPI Programmer's Guide* — An introduction to PatchAPI is stored in the **PatchAPI.pdf** file.
- *ProcControlAPI Programmer's Guide* — A detailed description of ProcControlAPI is stored in the **ProcControlAPI.pdf** file.
- *StackwalkerAPI Programmer's Guide* — A detailed description of StackwalkerAPI is stored in the **stackwalker.pdf** file.
- *SymtabAPI Programmer's Guide* — An introduction to SymtabAPI is stored in the **SymtabAPI.pdf** file.
- *InstructionAPI Reference Manual* — A detailed description of the InstructionAPI is stored in the **InstructionAPI.pdf** file.

For information on how to install this package on your system, see [Section 14.1, “Installing Dyninst”](#).

Online Documentation

- [Dyninst Home Page](#) — The project home page provides links to additional documentation and related publications.
- [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 7 SystemTap Beginners Guide](#) — The *SystemTap Beginners Guide* for Red Hat Enterprise Linux 7 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 7 SystemTap Tapset Reference](#) — The *SystemTap Tapset Reference* for Red Hat Enterprise Linux 7 provides further details about SystemTap.

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 11, SystemTap](#) provides an introduction to **SystemTap** and explains how to use it to monitor the activities of a running system.
- [Chapter 12, 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 13, OProfile](#) explains how to use **OProfile** to determine which sections of code consume the greatest amount of CPU time and why.

PART V. GETTING HELP

CHAPTER 15. ACCESSING RED HAT PRODUCT DOCUMENTATION

Red Hat Product Documentation located at <https://access.redhat.com/site/documentation/> serves as a central source of information. It is currently translated in 23 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 6.1 Release Notes](#) — The *Release Notes* for Red Hat Developer Toolset 6.1 contain more information.
- [Red Hat Software Collections Packaging Guide](#) — The *Software Collections Packaging Guide* explains the concept of Software Collections and documents how to create, build, and extend them.

RED HAT ENTERPRISE LINUX

- [Red Hat Enterprise Linux 6 Developer Guide](#) and [Red Hat Enterprise Linux 7 Developer Guide](#) — The *Developer Guides* for Red Hat Enterprise Linux 6 and 7 provide more 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 6 Installation Guide](#) and [Red Hat Enterprise Linux 7 Installation Guide](#) — The *Installation Guides* for Red Hat Enterprise Linux 6 and 7 explain 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 7 System Administrator's Guide](#) — The *System Administrator's Guide* for Red Hat Enterprise Linux 7 documents relevant information regarding the deployment, configuration, and administration of Red Hat Enterprise Linux 7.

CHAPTER 16. 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)**.

16.1. 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 workaround.

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 workaround 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 workaround.

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, see <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>.

16.2. ESCALATING AN ISSUE

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

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.

16.3. RE-OPENING A SERVICE REQUEST

If there is more relevant information regarding a closed service request (such as the problem reoccurring), you can re-open the request via the Red Hat Customer Portal at 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.

16.4. ADDITIONAL RESOURCES

For more information, see 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 the *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.

APPENDIX A. CHANGES IN VERSION 6.0

The sections below document features and compatibility changes introduced in Red Hat Developer Toolset 6.0.

A.1. CHANGES IN BINUTILS

Red Hat Developer Toolset 6.0 is distributed with **binutils 2.27**, which provides a number of bug fixes and feature enhancements over the version included in Red Hat Enterprise Linux and the previous version of Red Hat Developer Toolset. Below is a comprehensive list of new features in this release.

A.1.1. Changes Since Red Hat Enterprise Linux 6.8

The following features have been added since the release of **binutils** in Red Hat Enterprise Linux 6.8:

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

A.1.1.1. 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 6.0, the **gold** linker is not enabled by default. Users can explicitly switch between **ld** and **gold** by using the **alternatives** mechanism.

A.1.1.1.1. New Features

The following features have been added since the release of **binutils** included in Red Hat Enterprise Linux 6.8:

- 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.
- 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.
- A new **--compress-debug-sections** command line option has been added to enable the generation of compressed DWARF debug information sections in the relocatable output file.

A.1.1.1.2. Compatibility Changes

The following compatibility changes have been made since the release of **binutils** included in Red Hat Enterprise Linux 6.8:

- 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.

A.1.1.2. GNU Assembler

A.1.1.2.1. New Features

The following features have been added since the release of **binutils** included in Red Hat Enterprise Linux 6.8:

- The GNU Assembler no longer requires double ampersands in macros.
- A new **--compress-debug-sections** command line option has been added to enable the generation of compressed DWARF debug information sections in the relocatable output file.

Note that the setting of this option for the linker overrides the setting for the assembler. For example, if an object file contains a compressed debug section, but it is linked without **--compress-debug-sections** being passed to the linker, then those sections will be uncompressed as they are copied into the output binary.

- 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.

A.1.1.3. Other Binary Tools

A.1.1.3.1. New Features

The following features have been added since the release of **binutils** included in Red Hat Enterprise Linux 6.8:

- 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 **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.
- Support for reading the optimized debug information generated by the **dwz -m** tool has been added.
- The **devtoolset-2-binutils-devel** package now provides the **demangle.h** header file.

A.1.2. Changes Since Red Hat Developer Toolset 4.1

The following features have been added since the release of **binutils** in Red Hat Developer Toolset 4.1 and Red Hat Enterprise Linux 7.3:

- Support has been added for generating and using compressed debug sections. This reduces the size of binaries but can cause problems when other tools, not from the **binutils** package, try to examine them. For this reason, the actual generation of the compressed sections requires an explicit command-line option to be passed to the assembler or linker.

A.1.2.1. GNU Linker

A.1.2.1.1. New Features

The following features have been added since the release of **binutils** included in Red Hat Developer Toolset 4.1 and Red Hat Enterprise Linux 7.3:

- The linker now automatically enables the read-only run-time relocations (option **-z relro**) unless explicitly told otherwise. This helps to enhance the security of executables.
- The linker supports a new command-line option to define how orphan sections should be handled. Orphan sections are sections from input files whose placement in the output file is not defined by the linker script being used. The default behaviour is still maintained, but the option can be used to generate warning or error messages about them, or even to discard them entirely.
- The linker supports a new command line option, **--require-defined SYM**, which causes the linker to generate an error if *SYM* has not been defined by the time the end of the linking process has been reached.
- The linker supports a new command line option, **-z nodynamic-undefined-weak**, which stops it from generating dynamic relocations against undefined weak symbols in the executable being created.

A.1.2.2. GNU Assembler

A.1.2.2.1. New Features

The following features have been added since the release of **binutils** included in Red Hat Developer Toolset 4.1 and Red Hat Enterprise Linux 7.3:

- The assembler now supports the ARM v8.1 and ARM v8-M architectures, including the Adv.SIMD, LOR, PAN, Security, and DSP extensions.

- The assembler now allows ELF section flags and types to be set using numeric values as well as textual names.
- The assembler now accepts symbol and label names enclosed in double quotes ("), which allows them to contain characters that are not part of valid symbol names in high-level languages.

A.2. CHANGES IN ELFUTILS

Red Hat Developer Toolset 6.0 is distributed with **elfutils 0.167**, which provides a number of bug fixes and feature enhancements over the version included in the previous release of Red Hat Developer Toolset.

A.2.1. Changes Since Red Hat Developer Toolset 4.1

The following new ELF and DWARF string table creation functions have been added since the release of **elfutils** in Red Hat Developer Toolset 4.1: **dwelf_strtab_init**, **dwelf_strtab_add**, **dwelf_strtab_add_len**, **dwelf_strtab_finalize**, **dwelf_strent_off**, **dwelf_strent_str**, and **dwelf_strtab_free**.

A.3. CHANGES IN GCC

Red Hat Developer Toolset 6.0 is distributed with **GCC 6.2.1**, which provides a number of bug fixes and new features over the version included in Red Hat Enterprise Linux and the previous version of Red Hat Developer Toolset.

A.3.1. Changes Since Red Hat Developer Toolset 4.1

The following features have been added since the release of **GCC** in Red Hat Developer Toolset 4.1:

- The C++ compiler defaults to C++14 rather than C++98. Certain experimental C++17 language and runtime features have also been made available.
- Source locations are now tracked as ranges rather than points, which allows much richer diagnostics as well as "fixit" hints.
- New warnings have been implemented for statically detecting certain likely programming errors, including: negative shifts, shift overflow, tautological comparisons, null pointer dereferences, duplicated conditions, and misleading indentation.
- Various optimizer improvements have been added, particularly in alias analysis which helps to remove abstraction penalties for C++ code, improvements in the vectorizer, redundancy elimination, useless conditional elimination, and others.
- OpenMP 4.5 support for C and C++ has been added.
- Additional sanitizers to detect undefined behavior at runtime have been added.
- A new option, **-mfloat128**, has been implemented. It allows users to experiment with IEEE 128-bit floating point.

The 64-bit IBM POWER architecture now supports IEEE 128-bit floating-point using the **__float128** data type. Note that Red Hat Developer Toolset does not enable this support default, and **__float128** is not supported by the Red Hat Enterprise Linux 7 runtime.

Platform-Specific Improvements

The following improvements have been made to support for individual platforms:

- Support has been added for Intel Skylake processors with support for the AVX-512 extensions and the following instruction sets: Foundation (F), Byte and Word (BW), Doubleword and Quadword (DQ), Vector Length Extensions (VL), and Conflict Detection (CD).
- Support has been added for 64-bit ARM (AArch64) LSE extensions and support for new implementations and code-generation tuning for those implementations of the AArch64 ISA.
- Early support has been added for IEEE 128-bit floating point.
- Support for the z13 processor of the IBM z Systems architecture has been added.

A.4. CHANGES IN MAKE

Red Hat Developer Toolset 6.0 is distributed with **make 4.1**, which provides a number of bug fixes and enhancements over the version included in Red Hat Enterprise Linux.

A.4.1. Changes Since Red Hat Enterprise Linux 6.8

The supplied version of **make** supports integration with GNU **Guile** versions 1.8 as well as 2.0 and later.

The **make** version 4.1 released in Red Hat Developer Toolset 6.0 introduces the following compatibility changes:

- When the special target **.POSIX** is specified, **make** adheres to the POSIX requirements for handling backslash and newline. This changes how **make** handles backslash and newline in non-recipe lines:
 - Trailing spaces before the backslash are preserved
 - Each backslash and newline, including subsequent white space, is converted to single space
- The BSD **!=** shell assignment operator is now supported as an alternative to the **\$(shell ...)** function. As a consequence, variable names ending in exclamation mark, defined as **myvariable!= somevalue**, are now interpreted as shell assignments. To restore the previous behaviour, add white space after the intended variable name, such as **myvariable! = somevalue**.

A.5. CHANGES IN GDB

Red Hat Developer Toolset 6.0 is distributed with **GDB 7.12**, which provides a number of bug fixes and improvements over the Red Hat Enterprise Linux system version and the version included in the previous release of Red Hat Developer Toolset. Below is a comprehensive list of new features in this release.

A.5.1. Changes Since Red Hat Developer Toolset 4.1

The following features have been added since the release of **GDB** in Red Hat Developer Toolset 4.1:

New Features

- Support for tracepoints and fast tracepoints has been added in **GDBserver** for the following architectures:

- IBM z Systems
- IBM System z9
- IBM POWER

The support includes JIT compiling of conditional expressions in bytecode of fast tracepoints into native code.

- Support has been added for running interpreters on specified input and output devices.

GDB now supports a new mechanism that allows frontends to provide fully-featured **GDB** console views as a better alternative to building such views on top of the **-interpreter-exec console** command. See also the new **new-ui** command below. With that command, frontends can now start **GDB** in the traditional command-line mode running in an embedded terminal-emulator widget and create a separate MI interpreter running on a specified I/O device. In this way, **GDB** handles line editing, history, tab completion, and other tasks in the console all by itself, and the GUI uses the separate MI interpreter for its own control and synchronization, invisible to the command line.

- Support has been added for Fortran pointers to dynamic types.
- Support has been added for Fortran structures with fields of dynamic types and arrays of dynamic types.
- Support has been added to **GDBserver** for recording btrace without having to maintain an active **GDB** connection.
- Support has been added for a negative repeat count in the **x** command. This allows for examining memory backward from the given address. For example:

```
(gdb) bt
#0  Func1 (n=42, p=0x40061c "hohogoge") at main.cpp:4
#1  0x400580 in main (argc=1, argv=0x7fffffff5c8) at main.cpp:8
(gdb) x/-5i 0x0000000000400580
0x40056a <main(int, char**) +8>:      mov    %edi, -0x4(%rbp)
0x40056d <main(int, char**) +11>:     mov    %rsi, -0x10(%rbp)
0x400571 <main(int, char**) +15>:     mov    $0x40061c, %esi
0x400576 <main(int, char**) +20>:    mov    $0x2a, %edi
0x40057b <main(int, char**) +25>:
    callq 0x400536 <Func1(int, char const*)>
```

- Support has been added for multibit bitfields and enums in the target register descriptions.
- A new convenience function, **\$_as_string(val)**, based on Python, which returns the textual representation of a value, has been added. The function is useful for obtaining the text label of an enum value.
- Segmentation faults caused by Intel MPX boundary violations now display the type of violation (upper or lower), the memory address accessed, and the memory bounds, along with the signal received and code location. For example:

```
Program received signal SIGSEGV, Segmentation fault
Upper bound violation while accessing address 0x7fffffff3b3
Bounds: [lower = 0x7fffffff390, upper = 0x7fffffff3a3]
0x0000000000400d7c in upper () at i386-mpx-sigsegv.c:68
```

New Commands

The following new commands have been added:

skip -file *file*, skip -gfile *file-glob-pattern*, skip -function *function*, skip -rfunction *regular-expression*

A generalized form of the **skip** command, with new support for glob-style file names and regular expressions for function names. Additionally, a file spec and a function spec may now be combined.

maint info line-table *REGEXP*

Show the contents of the internal line-table data structure.

maint selftest

Run all compiled in **GDB** unit tests.

new-ui *INTERP TTY*

Start a new user interface instance running *INTERP* as interpreter, using the *TTY* file for input and output.

The following command has been enhanced:

- When issued with the **group:** or **g:** prefix, the **catch syscall** command now supports catching groups of related system calls.

Python Scripting Support

Python scripting support has been improved:

- A new attribute, **pending**, has been added to **`gdb.Breakpoint`** objects, which indicates whether the breakpoint is pending.
- Three new events related to breakpoints have been added: **`gdb.breakpoint_created`**, **`gdb.breakpoint_modified`**, and **`gdb.breakpoint_deleted`**.

Change in the Machine Interface Interpreter (GDB/MI)

The following change has been made the GDB/MI.

- The **`=record-started`** async record now includes the method and format used for recording. For example:

```
| =record-started, thread-group="i1", method="btrace", format="bts"
```

- The **`=thread-selected`** async record now includes the frame field. For example:

```
| =thread-selected, id="3", frame={level="0", addr="0x00000000004007c0"}
```

A.6. CHANGES IN OPROFILE

Red Hat Developer Toolset 6.0 is distributed with **OProfile 1.1.0**, which provides a number of bug fixes and feature enhancements over the version included in the previous release of Red Hat Developer Toolset. Below is a list of new features in this release.

A.6.1. Changes Since Red Hat Developer Toolset 4.1

The following features have been added since the release of **OProfile** included in Red Hat Developer Toolset 4.1:

- Support has been added for the following processors:
 - IBM z13 (only the **ocount** is supported)
 - Applied Micro X-Gene
- Event lists for IBM POWER8 processors have been updated.
- Support has been added for the following Intel processors:
 - 2nd Generation Intel Atom (Goldmont)
 - 6th Generation Intel Core, CPUID model 0x55 (Skylake)
 - 7th Generation Intel Core (Kaby Lake)
- The search logic has been corrected to properly store data in an archive and make use of the archive data.
- **operf** now only starts the command to be measured if the performance monitoring hardware is properly set up. This fixes a bug that caused **operf** to start the command that is to be measured regardless of whether the performance monitoring set up was successful.

A.7. CHANGES IN STRACE

Red Hat Developer Toolset 6.0 is distributed with **strace 4.12**, which provides a number of bug fixes.

A.8. CHANGES IN SYSTEMTAP

Red Hat Developer Toolset 6.0 is distributed with **SystemTap 3.0**, which provides a number of bug fixes and feature enhancements over the version included in the previous release of Red Hat Developer Toolset. Below is a list of new features in this release.

A.8.1. Changes Since Red Hat Developer Toolset 4.1

The following main features have been added since the release of **SystemTap** included in Red Hat Developer Toolset 4.1:

- Experimental monitor and interactive modes.
- Optimized associative arrays.
- Two types of function overloading.
- Probe point brace-pattern expansion.
- Security band-aid samples.
- Improved string quoting and escaping.
- Pretty printing of array aggregates.

- Private scoping for variables.



NOTE

Incompatibility problems with old scripts can be resolved using the backward-compatibility option, `--compatible version`, where *version* is the version of **SystemTap** for which the script was written.

A.9. CHANGES IN DYNINST

Red Hat Developer Toolset 6.0 is distributed with **dyninst 9.2.0**, which provides a number of bug fixes and feature enhancements over the version included in the previous release of Red Hat Developer Toolset. Below is a list of new features in this release.

A.9.1. Changes Since Red Hat Developer Toolset 4.1

The following features have been added since the release of **dyninst** included in Red Hat Developer Toolset 4.1:

- Expanded instruction support for x86 architectures to include the AVX, AVX2, and AVX-512 extensions, and many other fixes in x86 decoding.
- Fixed rewriting for position-independent executables (PIE).
- Improved **SymtabAPI** coverage using DWARF information.
- Improved analysis of jump tables.
- Safety fixes for memory allocation and tramp guards in the **RTlib** library.

A.10. CHANGES IN VALGRIND

Red Hat Developer Toolset 6.0 is distributed with **Valgrind 3.12.0**, which provides a number of bug fixes and enhancements over the version included in Red Hat Enterprise Linux and the previous version of Red Hat Developer Toolset.

A.10.1. Changes Since Red Hat Developer Toolset 4.1

The following bug fixes and enhancements have been added since the release of **Valgrind** in Red Hat Developer Toolset 4.1:

- Added meta mempool support for describing a custom allocator. This feature:
 - Automatically frees all chunks assuming that destroying a pool destroys all objects in the pool.
 - Uses itself to allocate other memory blocks.
- The maximum number of callers in a suppression entry is now equal to the maximum size for the `--num-callers` option (500). Note that setting the `--gen-suppressions=yes|all` option similarly generates suppression containing up to `--num-callers` frames.
- The **gdbserver** provided by **Valgrind** now accepts the `catch syscall` command.

- JIT has been improved to lower the cost of instrumenting code blocks for the most common use case (AMD64 and Intel 64 with **Memcheck**). The cost has been reduced by 10-15%.
- Limited support has been added for certain AMD64 FMA4 instructions.
- Support for the following architectures has been added:
 - 64-bit ARM (AArch64)
 - IBM System z9

The valgrind-openmpi package has been removed from this release.

APPENDIX B. CHANGES IN VERSION 6.1

The sections below document features and compatibility changes introduced in Red Hat Developer Toolset 6.1.

B.1. CHANGES IN ELFUTILS

Red Hat Developer Toolset 6.1 is distributed with **elfutils 0.168**, which provides a number of bug fixes and feature enhancements over the version included in the previous release of Red Hat Developer Toolset.

B.1.1. Changes Since Red Hat Developer Toolset 6.0

The following features have been added since the release of **elfutils** included in Red Hat Developer Toolset 6.0:

- The **eu-readelf** utility now allows specifying the symbol table section by name with parameter **-symbols[=SECTION]**. This enables showing just the symbol table of the specified ELF section name instead of symbols from all symbol tables.
- Return type of functions **gelf_newehdr** and **gelf_newehdr** in the **libelf** library changed from **unsigned long int** to **void***. The new signatures are abi-compatible on all supported architectures. This change increases compatibility of code with implementations of **libelf** on other platforms.
- The **DW_LANG_PL1** constant in **dwarf.h** header file from **libdw** has been changed to **DW_LANG_PLI**. **libdw** still includes the old name as an alias.

The version of **elfutils** included in Red Hat Developer Toolset 6.1 includes one bug fix:

- The **devtoolset-6-elfutils-libs** package provided the same shared library object names as the base RHEL **elfutils-libs** package. Consequently, if a package required one of the provided **elfutils** libraries (**libelf**, **libdw**, or **libasm**), **devtoolset-6-elfutils-libs** was installed instead of the base Red Hat Enterprise Linux **elfutils-libs** package. With this update, **devtoolset-6-elfutils-libs** no longer provides the same library names as the system **elfutils-libs** package. As a result, when a package needs an **elfutils-libs** library, the system **elfutils-libs** package is correctly installed. (BZ#[1412305](#))

B.2. CHANGES IN GCC

Red Hat Developer Toolset 6.1 is distributed with **GCC 6.3.1**, which provides numerous bug fixes over the version included in Red Hat Enterprise Linux and the previous version of Red Hat Developer Toolset.

B.3. CHANGES IN GDB

Red Hat Developer Toolset 6.1 is distributed with **GDB 7.12.1**, which provides a number of bug fixes over the version included in the previous release of Red Hat Developer Toolset.

B.4. CHANGES IN DEVTOOLSET-6-TOOLCHAIN-RHEL7

Red Hat Developer Toolset 6.1 is distributed with **devtoolset-6-toolchain-rhel7**, a pre-built docker image that includes other Red Hat Developer Toolset components.

B.4.1. Changes Since Red Hat Developer Toolset 6.0

For the changes since the release of **devtoolset-6-toolchain-rhel7** included in Red Hat Developer Toolset 6.0, see [Section B.2, “Changes in GCC”](#).

B.5. CHANGES IN DEVTOOLSET-6-PERFTOOLS-RHEL7

Red Hat Developer Toolset 6.1 is distributed with **devtoolset-6-perftools-rhel7**, a pre-built docker image that includes other Red Hat Developer Toolset components.

B.5.1. Changes Since Red Hat Developer Toolset 6.0

For the changes since the release of **devtoolset-6-perftools-rhel7** included in Red Hat Developer Toolset 6.0, see [Section B.1, “Changes in elfutils”](#).

APPENDIX C. REVISION HISTORY

Revision 6.1-3 Clarified C++ compatibility information.	Fri 26 May 2017	Vladimír Slávik
Revision 6.1-2 Release of Red Hat Developer Toolset 6.1 User Guide.	Wed 26 Apr 2017	Vladimír Slávik
Revision 6.1-1 Release of Red Hat Developer Toolset 6.1 Beta User Guide.	Wed 05 Apr 2017	Vladimír Slávik
Revision 6.0-9 Removed outdated information.	Tue 04 Apr 2017	Lenka Špačková
Revision 6.0-8 Release with minor fixes.	Fri 31 Mar 2017	Vladimír Slávik
Revision 6.0-7 Release of Red Hat Developer Toolset 6.0 User Guide.	Sun 13 Nov 2016	Robert Krátký
Revision 0.0-36 Release of Red Hat Developer Toolset 4.1 User Guide.	Tue 24 May 2016	Robert Krátký
Revision 0.0-33 Release of Red Hat Developer Toolset 4.0 User Guide.	Fri 13 Nov 2015	Robert Krátký
Revision 0.0-29 Re-release of Red Hat Developer Toolset 4.0 Beta User Guide with a section on docker-formatted container images.	Wed 4 Nov 2015	Robert Krátký
Revision 0.0-28 Release of Red Hat Developer Toolset 4.0 Beta User Guide.	Wed 14 Oct 2015	Robert Krátký
Revision 0.0-25 Update to reflect RHSCL 2.0 GA.	Thu 4 June 2015	Robert Krátký
Revision 0.0-23 Release of Red Hat Developer Toolset 3.1 User Guide.	Thu 23 Apr 2015	Robert Krátký
Revision 0.0-17 Release of Red Hat Developer Toolset 3.1 Beta User Guide.	Tue 10 Mar 2015	Robert Krátký
Revision 0.0-16 Release of Red Hat Developer Toolset 3.0 User Guide with minor post-GA fixes.	Thu 13 Nov 2014	Robert Krátký
Revision 0.0-14 Release of Red Hat Developer Toolset 3.0 User Guide.	Thu 30 Oct 2014	Robert Krátký
Revision 0.0-10 Release of Red Hat Developer Toolset 3.0 Beta-2 User Guide.	Tue 07 Oct 2014	Robert Krátký
Revision 0.0-8 Release of Red Hat Developer Toolset 3.0 Beta-1 User Guide.	Tue Sep 09 2014	Robert Krátký

INDEX

A

addr2line

features, [New Features](#)

overview, [binutils](#)

usage, [Using Other Binary Tools](#)

ar

overview, [binutils](#)

usage, [Using Other Binary Tools](#)

as (see [GNU assembler](#))

assembling (see [GNU assembler](#))

B

binutils

documentation, [Additional Resources](#)

installation, [Installing binutils](#)

overview, [binutils](#)

usage, [Using the GNU Assembler](#), [Using the GNU Linker](#), [Using Other Binary Tools](#)

version, [About Red Hat Developer Toolset](#), [binutils](#)

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

compatibility, [C++ Compatibility](#)

compiling, [Using the C++ Compiler](#), [Preparing a Program for Debugging](#)

running, [Running a C++ Program](#)

support, [GNU C++ Compiler](#)

c++filt

overview, [binutils](#)

usage, [Using Other Binary Tools](#)

Cachegrind

overview, [Valgrind](#)

usage, [Using Valgrind](#)

Callgrind

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

compatibility

Red Hat Developer Toolset, [Compatibility](#)

compiling (see GNU Compiler Collection)

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](#)

dwp

overview, [binutils](#)

usage, [Using Other Binary Tools](#)

dwz

documentation, [Additional Resources](#)

installation, [Installing dwz](#)

overview, [dwz](#)

usage, [Using dwz](#)

version, [About Red Hat Developer Toolset, dwz](#)

Dyninst

documentation, [Additional Resources](#)

installation, [Installing Dyninst](#)

overview, [Dyninst](#)

usage, [Using Dyninst](#)

version, [About Red Hat Developer Toolset, Dyninst](#)

E

elfedit

features, [New Features](#)

overview, [binutils](#)

usage, [Using Other Binary Tools](#)

elfutils

documentation, [Additional Resources](#)

installation, [Installing elfutils](#)

overview, [elfutils](#)

usage, [Using elfutils](#)

version, [About Red Hat Developer Toolset, elfutils](#)

eu-addr2line

overview, [elfutils](#)

usage, [Using elfutils](#)

eu-ar

overview, [elfutils](#)

usage, [Using elfutils](#)

eu-elfcmp

overview, [elfutils](#)

usage, [Using elfutils](#)

eu-elflint

overview, [elfutils](#)

usage, [Using elfutils](#)

eu-findtextrel

overview, [elfutils](#)

usage, [Using elfutils](#)

eu-make-debug-archive

overview, [elfutils](#)

usage, [Using elfutils](#)

eu-nm

overview, [elfutils](#)

usage, [Using elfutils](#)

eu-objdump

overview, [elfutils](#)

usage, [Using elfutils](#)

eu-ranlib

overview, [elfutils](#)

usage, [Using elfutils](#)

eu-readelf

overview, [elfutils](#)

usage, [Using elfutils](#)

eu-size

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

eu-stack
overview, [elfutils](#)

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

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

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

F

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
documentation, [Additional Resources](#)
installation, [Installing binutils](#)
overview, [binutils](#)
usage, [Using the GNU Assembler](#)

GNU Binutils (see binutils)
GNU Compiler Collection
C support, [GNU C Compiler](#)

C++ support, [GNU C++ Compiler](#)

documentation, [Additional Resources](#)

features, [Main 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, [About Red Hat Developer Toolset](#), [GNU Compiler Collection \(GCC\)](#)

GNU Debugger

documentation, [Additional Resources](#)

features, [Main 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, [About Red Hat Developer Toolset](#), [GNU Debugger \(GDB\)](#)

GNU linker

documentation, [Additional Resources](#)

installation, [Installing binutils](#)

overview, [binutils](#)

usage, [Using the GNU Linker](#)

GNU make

documentation, [Additional Resources](#)

installation, [Installing make](#)

version, [GNU make](#)

gprof

overview, [binutils](#)

usage, [Using Other Binary Tools](#)

GSS (see Global Support Services)

H

Helgrind

overview, [Valgrind](#)

usage, [Using Valgrind](#)

help

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

L

ld (see GNU linker)

linking (see GNU linker)

ltrace

documentation, [Additional Resources](#)

installation, [Installing ltrace](#)

overview, [ltrace](#)

usage, [Using ltrace](#)

version, [About Red Hat Developer Toolset](#), [ltrace](#)

M

make

building, [Using make](#)

overview, [GNU make](#)

version, [About Red Hat Developer Toolset](#)

Makefile

usage, [Using Makefiles](#)

Massif

overview, [Valgrind](#)

usage, [Using Valgrind](#)

Memcheck

overview, [Valgrind](#)

usage, [Using Valgrind](#)

memstomp

documentation, [Additional Resources](#)

installation, [Installing memstomp](#)

overview, [memstomp](#)

usage, [Using memstomp](#)

version, [About Red Hat Developer Toolset](#)

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](#)

ocount

overview, [OProfile](#)

opannotate

overview, [OProfile](#)

usage, [Using OProfile](#)

oparchive

overview, [OProfile](#)

usage, [Using OProfile](#)

operf

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

documentation, [Additional Resources](#)
installation, [Installing OProfile](#)
overview, [OProfile](#)
usage, [Using OProfile](#)
version, [About Red Hat Developer Toolset](#), [OProfile](#)

oprofiled

overview, [OProfile](#)

R

ranlib

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

readelf

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

Red Hat Developer Toolset

compatibility, [Compatibility](#)
Container Images, [Using Red Hat Developer Toolset Container Images](#)
Docker, [Using Red Hat Developer Toolset Container Images](#)
Docker-formatted container images, [Using Red Hat Developer Toolset Container Images](#)
Dockerfiles, [Using Red Hat Developer Toolset Container Images](#)
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](#), [Using Dyninst with SystemTap](#)

stap-merge

overview, [SystemTap](#)

usage, [Using SystemTap](#)

stap-prep

overview, [SystemTap](#)

usage, [Installing SystemTap](#)

stap-report

overview, [SystemTap](#)

usage, [Using SystemTap](#)

stap-server

overview, [SystemTap](#)

stapdyn

overview, [SystemTap](#)

staprun

overview, [SystemTap](#)

usage, [Using SystemTap](#)

stapsh

overview, [SystemTap](#)

usage, [Using SystemTap](#)

strace

documentation, [Additional Resources](#)

installation, [Installing strace](#)

overview, [strace](#)

usage, [Using strace](#)

version, [About Red Hat Developer Toolset](#), [strace](#)

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

documentation, [Additional Resources](#)

installation, [Installing SystemTap](#)

overview, [SystemTap](#)

usage, [Using SystemTap](#), [Using Dyninst with SystemTap](#)

version, [About Red Hat Developer Toolset](#), [SystemTap](#)

V

Valgrind

documentation, [Additional Resources](#)

installation, [Installing Valgrind](#)

overview, [Valgrind](#)

usage, [Using Valgrind](#)

version, [About Red Hat Developer Toolset](#), [Valgrind](#)

version

version, [memstomp](#)