



Red Hat Developer Toolset 10

User Guide

Installing and Using Red Hat Developer Toolset

Red Hat Developer Toolset 10 User Guide

Installing and Using Red Hat Developer Toolset

Olga Tikhomirova
Red Hat Customer Content Services
otikhomi@redhat.com

Zuzana Zoubková
Red Hat Customer Content Services

Jaromír Hradílek
Red Hat Customer Content Services

Matt Newsome
Red Hat Software Engineering

Robert Krátký
Red Hat Customer Content Services

Vladimír Slávik
Red Hat Customer Content Services

Legal Notice

Copyright © 2021 Red Hat, Inc.

The text of and illustrations in this document are licensed by Red Hat under a Creative Commons Attribution–Share Alike 3.0 Unported license ("CC-BY-SA"). An explanation of CC-BY-SA is available at

<http://creativecommons.org/licenses/by-sa/3.0/>

. In accordance with CC-BY-SA, if you distribute this document or an adaptation of it, you must provide the URL for the original version.

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

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

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

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

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

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

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

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

All other trademarks are the property of their respective owners.

Abstract

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

MAKING OPEN SOURCE MORE INCLUSIVE	7
PART I. INTRODUCTION	8
CHAPTER 1. RED HAT DEVELOPER TOOLSET	9
1.1. ABOUT RED HAT DEVELOPER TOOLSET	9
What Is New in Red Hat Developer Toolset 10.1	9
1.2. MAIN FEATURES	10
1.3. COMPATIBILITY	11
Architecture support	11
1.4. GETTING ACCESS TO RED HAT DEVELOPER TOOLSET	11
1.4.1. Using Red Hat Software Collections	11
1.5. INSTALLING RED HAT DEVELOPER TOOLSET	13
1.5.1. Installing All Available Components	13
1.5.2. Installing Individual Package Groups	13
1.5.3. Installing Optional Packages	14
1.5.4. Installing Debugging Information	14
1.6. UPDATING RED HAT DEVELOPER TOOLSET	15
1.6.1. Updating to a Minor Version	15
1.6.2. Updating to a Major Version	15
1.7. UNINSTALLING RED HAT DEVELOPER TOOLSET	15
1.8. USING RED HAT DEVELOPER TOOLSET CONTAINER IMAGES	16
1.9. ADDITIONAL RESOURCES	16
Online Documentation	16
See Also	17
PART II. DEVELOPMENT TOOLS	18
CHAPTER 2. GNU COMPILER COLLECTION (GCC)	19
2.1. GNU C COMPILER	19
2.1.1. Installing the C Compiler	19
2.1.2. Using the C Compiler	19
2.1.3. Running a C Program	20
2.2. GNU C++ COMPILER	20
2.2.1. Installing the C++ Compiler	21
2.2.2. Using the C++ Compiler	21
2.2.3. Running a C++ Program	22
2.2.4. C++ Compatibility	22
2.2.4.1. C++ ABI	23
2.3. GNU FORTRAN COMPILER	23
2.3.1. Installing the Fortran Compiler	24
2.3.2. Using the Fortran Compiler	24
2.3.3. Running a Fortran Program	25
2.4. SPECIFICS OF GCC IN RED HAT DEVELOPER TOOLSET	25
2.5. ADDITIONAL RESOURCES	26
Installed Documentation	26
Online Documentation	26
See Also	26
CHAPTER 3. GNU MAKE	28
3.1. INSTALLING MAKE	28
3.2. USING MAKE	28

3.3. USING MAKEFILES	29
3.4. ADDITIONAL RESOURCES	30
Installed Documentation	30
Online Documentation	30
See Also	30
CHAPTER 4. BINUTILS	32
4.1. INSTALLING BINUTILS	33
4.2. USING THE GNU ASSEMBLER	33
4.3. USING THE GNU LINKER	33
4.4. USING OTHER BINARY TOOLS	34
4.5. SPECIFICS OF BINUTILS IN RED HAT DEVELOPER TOOLSET	35
4.6. ADDITIONAL RESOURCES	35
Installed Documentation	35
Online Documentation	35
See Also	35
CHAPTER 5. ELFUTILS	37
5.1. INSTALLING ELFUTILS	38
5.2. USING ELFUTILS	38
5.3. ADDITIONAL RESOURCES	38
See Also	38
CHAPTER 6. DWZ	39
6.1. INSTALLING DWZ	39
6.2. USING DWZ	39
6.3. ADDITIONAL RESOURCES	39
Installed Documentation	39
See Also	39
CHAPTER 7. ANNOBIN	41
7.1. INSTALLING ANNOBIN	41
7.2. USING ANNOBIN PLUGIN	41
7.3. USING ANNOCHECK	41
7.4. ADDITIONAL RESOURCES	42
Installed Documentation	42
PART III. DEBUGGING TOOLS	43
CHAPTER 8. GNU DEBUGGER (GDB)	44
8.1. INSTALLING THE GNU DEBUGGER	44
8.2. PREPARING A PROGRAM FOR DEBUGGING	44
Compiling Programs with Debugging Information	44
Installing Debugging Information for Existing Packages	45
8.3. RUNNING THE GNU DEBUGGER	45
8.4. LISTING SOURCE CODE	46
8.5. SETTING BREAKPOINTS	47
Setting a New Breakpoint	47
Listing Breakpoints	48
Deleting Existing Breakpoints	48
8.6. STARTING EXECUTION	49
8.7. DISPLAYING CURRENT VALUES	49
8.8. CONTINUING EXECUTION	50
8.9. ADDITIONAL RESOURCES	50
Installed Documentation	50

Online Documentation	51
See Also	51
CHAPTER 9. STRACE	52
9.1. INSTALLING STRACE	52
9.2. USING STRACE	52
9.2.1. Redirecting Output to a File	52
9.2.2. Tracing Selected System Calls	53
9.2.3. Displaying Time Stamps	54
9.2.4. Displaying a Summary	55
9.2.5. Tampering with System Call Results	55
9.3. ADDITIONAL RESOURCES	56
Installed Documentation	56
See Also	56
CHAPTER 10. LTRACE	57
10.1. INSTALLING LTRACE	57
10.2. USING LTRACE	57
10.2.1. Redirecting Output to a File	57
10.2.2. Tracing Selected Library Calls	58
10.2.3. Displaying Time Stamps	59
10.2.4. Displaying a Summary	59
10.3. ADDITIONAL RESOURCES	60
Installed Documentation	60
Online Documentation	60
See Also	60
CHAPTER 11. MEMSTOMP	61
11.1. INSTALLING MEMSTOMP	62
11.2. USING MEMSTOMP	62
11.3. ADDITIONAL RESOURCES	64
Installed Documentation	64
See Also	64
PART IV. PERFORMANCE MONITORING TOOLS	65
CHAPTER 12. SYSTEMTAP	66
12.1. INSTALLING SYSTEMTAP	66
12.2. USING SYSTEMTAP	67
12.3. ADDITIONAL RESOURCES	67
Installed Documentation	67
Online Documentation	67
See Also	68
CHAPTER 13. VALGRIND	69
13.1. INSTALLING VALGRIND	69
13.2. USING VALGRIND	70
13.3. ADDITIONAL RESOURCES	70
Installed Documentation	70
Online Documentation	70
See Also	71
CHAPTER 14. OPROFILE	72
14.1. INSTALLING OPROFILE	72
14.2. USING OPROFILE	72

14.3. ADDITIONAL RESOURCES	73
Installed Documentation	73
Online Documentation	73
See Also	74
CHAPTER 15. DYNINST	75
15.1. INSTALLING DYNINST	75
15.2. USING DYNINST	75
15.2.1. Using Dyninst with SystemTap	75
15.2.2. Using Dyninst as a Stand-alone Library	76
15.3. ADDITIONAL RESOURCES	80
Installed Documentation	80
Online Documentation	80
See Also	80
PART V. COMPILER TOOLSETS	82
CHAPTER 16. COMPILER TOOLSETS DOCUMENTATION	83
PART VI. GETTING HELP	84
CHAPTER 17. ACCESSING RED HAT PRODUCT DOCUMENTATION	85
Red Hat Developer Toolset	85
Red Hat Enterprise Linux	85
CHAPTER 18. CONTACTING GLOBAL SUPPORT SERVICES	86
18.1. GATHERING REQUIRED INFORMATION	86
Background Information	86
Diagnostics	86
Account and Contact Information	86
Issue Severity	87
18.2. ESCALATING AN ISSUE	87
18.3. RE-OPENING A SERVICE REQUEST	87
18.4. ADDITIONAL RESOURCES	88
Online Documentation	88
APPENDIX A. CHANGES IN VERSION 10.0	89
A.1. CHANGES IN GCC	89
General improvements	89
Language-specific improvements	90
Target-specific improvements	93
A.2. CHANGES IN BINUTILS	93
The assembler	93
The linker	93
Other binary utilities	94
A.3. CHANGES IN ELFUTILS	95
A.4. CHANGES IN GDB	95
New features	95
New convenience variables and functions	95
New and improved commands	96
New commands	96
Changed commands	96
Settings	97
Language-specific improvements	99
Python API	99

Machine Interface (MI) improvements	99
A.5. CHANGES IN STRACE	100
Changes in behavior	100
Improvements	100
Bug fixes	102
A.6. CHANGES IN SYSTEMTAP	102
A.7. CHANGES IN VALGRIND	102
APPENDIX B. CHANGES IN VERSION 10.1	104
B.1. CHANGES IN GCC	104
B.2. CHANGES IN ELFUTILS	104
B.3. CHANGES IN STRACE	104
B.4. CHANGES IN DYNINST	105
B.5. CHANGES IN SYSTEMTAP	105

MAKING OPEN SOURCE MORE INCLUSIVE

Red Hat is committed to replacing problematic language in our code, documentation, and web properties. We are beginning with these four terms: master, slave, blacklist, and whitelist. Because of the enormity of this endeavor, these changes will be implemented gradually over several upcoming releases. For more details, see [our CTO Chris Wright's message](#).

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 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 10.1

Since Red Hat Developer Toolset 4.1, the Red Hat Developer Toolset content is also available in the ISO format together with the rest of Red Hat Software Collections content 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	10.2.1	A portable compiler suite with support for C, C++, and Fortran.
binutils	2.35	A collection of binary tools and other utilities to inspect and manipulate object files and binaries.
elfutils	0.182	A collection of binary tools and other utilities to inspect and manipulate ELF files.
dwz	0.12	A tool to optimize DWARF debugging information contained in ELF shared libraries and ELF executables for size.
GDB	9.2	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	5.7	A debugging tool to monitor system calls that a program uses and signals it receives.
memstomp	0.15	A debugging tool to identify calls to library functions with overlapping memory regions that are not allowed by various standards.
SystemTap	4.4	A tracing and probing tool to monitor the activities of the entire system without the need to instrument, recompile, install, and reboot.

Name	Version	Description
Valgrind	3.16.1	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	1.4.0	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	10.2.1	A library for instrumenting and working with user-space executables during their execution.
make	4.2.1	A dependency-tracking build automation tool.
annobin	9.23	A build security checking 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 `scl` utility.

1.2. MAIN FEATURES

Red Hat Developer Toolset 10.1 brings the following changes:

- The Red Hat Developer Toolset version of the **GNU Compiler Collection (GCC)** has been upgraded to version 10.2.1 with many new features and bug fixes.
- The Red Hat Developer Toolset version of the **GNU Debugger (GDB)** has been upgraded to version 9.2 with many new features and bug fixes.

For a full list of changes and features introduced in this release, see [Appendix B, Changes in Version 10.1](#).

1.3. COMPATIBILITY

Red Hat Developer Toolset 10.1 is available for Red Hat Enterprise Linux 7 for a number of architectures.

For ABI compatibility information, see [Section 2.2.4, "C++ Compatibility"](#).

Table 1.2. Red Hat Developer Toolset 10.1 Compatibility

	Runs on Red Hat Enterprise Linux 7.6	Runs on Red Hat Enterprise Linux 7.7	Runs on Red Hat Enterprise Linux 7.9
Built with Red Hat Enterprise Linux 7.6	Supported	Supported	Supported
Built with Red Hat Enterprise Linux 7.7	Not supported	Supported	Supported
Built with Red Hat Enterprise Linux 7.9	Not Supported	Not Supported	Supported

Architecture support

Red Hat Developer Toolset is available on the following architectures:

- The 64-bit Intel and AMD architectures
- IBM Power Systems, big endian
- IBM Power Systems, little endian
- 64-bit IBM Z

1.4. GETTING ACCESS TO RED HAT DEVELOPER TOOLSET

Red Hat Developer Toolset is an offering distributed as a part of Red Hat Software Collections.

This content set is available to customers with Red Hat Enterprise Linux 7 subscriptions listed at <https://access.redhat.com/solutions/472793>.

Enable Red Hat Developer Toolset by using Red Hat Subscription Management. For information on how to register your system with this subscription management service, see the [Red Hat Subscription Management](#) collection of guides.

1.4.1. Using Red Hat Software Collections

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, 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:

```
# 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:

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

3. Determine the exact name of the Red Hat Software Collections repository. 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-rhscl-version-rpms  
rhel-variant-rhscl-version-debug-rpms  
rhel-variant-rhscl-version-source-rpms
```

In addition, certain packages, such as **devtoolset-10-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 (**7**).

4. Enable the repositories from step no. 3:


```
# 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.

1.5. 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-10** package:

```
# yum install devtoolset-10
```

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 1.3, “Red Hat Developer Toolset Meta Packages”](#).

Table 1.3. Red Hat Developer Toolset Meta Packages

Package Name	Description	Installed Components
devtoolset-10-perftools	Performance monitoring tools	SystemTap, Valgrind, OProfile, Dyninst

Package Name	Description	Installed Components
devtoolset-10-toolchain	Development and debugging tools	GCC, make, GDB, binutils, elfutils, dwz, memstomp, strace, ltrace

To install any of these meta packages:

```
# 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:

```
# yum install devtoolset-10-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:

```
$ yum list available devtoolset-10-*
```

To install any of these optional packages:

```
# 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-10-gdb-gdbserver** and **devtoolset-10-gdb-doc** packages:

```
# yum install devtoolset-10-gdb-gdbserver devtoolset-10-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:


```
# debuginfo-install package_name
```

For example, to install debugging information for the **devtoolset-10-dwz** package:

```
# debuginfo-install devtoolset-10-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, “Getting Access to Red Hat](#)

Developer Toolset”. For more information on how to get access to debuginfo packages, see <https://access.redhat.com/site/solutions/9907>.



NOTE

The `devtoolset-10-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 10.1 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, 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 10.1 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:

```
# yum remove devtoolset-10* 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 10.1 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.

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 10.1 Release Notes](#) – The *Release Notes* for Red Hat Developer Toolset 10.1 contain more information.
- [Red Hat Enterprise Linux 7 Developer Guide](#) – The *Developer Guide* for Red Hat Enterprise Linux 7 provides more information on the **Eclipse** IDE, libraries and runtime support, compiling and building, debugging, and profiling on these systems.
- [Red Hat Enterprise Linux 7 Installation Guide](#) – The *Installation Guide* for Red Hat Enterprise Linux 7 explains how to obtain, install, and update the system.
- [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.

- [Getting Started with Containers](#) – The guide contains a comprehensive overview of information about building and using container images on Red Hat Enterprise Linux 7 and Red Hat Enterprise Linux Atomic Host.

See Also

- [Appendix B, Changes in Version 10.1](#) – A list of changes and improvements over the version of the Red Hat Developer Toolset tools 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 10.2.1**. 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-10-gcc** package and is automatically installed with **devtoolset-10-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:

```
$ scl enable devtoolset-10 '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,;


```
$ scl enable devtoolset-10 '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:

```
$ scl enable devtoolset-10 'gcc -o output_file object_file...'
```

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

```
$ scl enable devtoolset-10 'bash'
```



NOTE

To verify the version of **gcc** you are using at any point:

```
$ 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
```

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;
}
```

Compile this source code on the command line by using the **gcc** compiler from Red Hat Developer Toolset:

```
$ scl enable devtoolset-10 '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 run it:

```
$/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-10-gcc-c++` package and is automatically installed with the `devtoolset-10-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:

```
$ scl enable devtoolset-10 '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:

```
$ scl enable devtoolset-10 '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:

```
$ scl enable devtoolset-10 'g++ -o output_file object_file..'
```

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

```
$ scl enable devtoolset-10 'bash'
```

NOTE

To verify the version of `g++` you are using at any point:

```
$ 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
```

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;
}
```

Compile this source code on the command line by using the **g++** compiler from Red Hat Developer Toolset:

```
$ scl enable devtoolset-10 '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 run it:

```
$/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:

```
$/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 to 10 in any **-std** mode are compatible with any other of those compilers in C++98 mode.

A compiler in C++11, C++14, or C++17 mode is only guaranteed to be compatible with another compiler in those same modes if they are from the same release series.

Supported examples:

- C++11 and C++11 from Red Hat Developer Toolset 6.x
- C++14 and C++14 from Red Hat Developer Toolset 6.x
- C++17 and C++17 from Red Hat Developer Toolset 10.x



IMPORTANT

- The GCC compiler in Red Hat Developer Toolset 10.x can build code using C++20 but this capability is experimental and not supported by Red Hat.
- 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 by the Red Hat Developer Toolset toolchain 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 GCC.

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 the C++14 language version is supported in Red Hat Developer Toolset when all C++ objects compiled with the respective flag have been built using Red Hat Developer Toolset 6 or later. Objects compiled by the system GCC in its default mode of C++98 are also compatible, but objects compiled with the system GCC in C++11 or C++14 mode are not compatible.

Starting with Red Hat Developer Toolset 10.x, using the C++17 language version is no longer experimental and is supported by Red Hat.



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 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 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 lengths. The compiler in Red Hat Developer Toolset uses the new mangling by default. It is possible to use the previous standard mangling by adding the **-fabi-version=2** or **-fabi-version=3** options to GCC C++ compiler calls. To display a warning about code that uses the old mangling, use the **-Wabi** option.

On Red Hat Enterprise Linux 7, the GCC C++ compiler still uses the old mangling by default, but emits aliases with the new mangling on targets that support strong aliases. It is possible to use the new standard mangling by adding the **-fabi-version=4** option to compiler calls. To display a warning about code that uses the old mangling, use the **-Wabi** option.

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-10-gcc-gfortran** package and is automatically installed with **devtoolset-10-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:

```
$ scl enable devtoolset-10 '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:


```
$ scl enable devtoolset-10 '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:

```
$ scl enable devtoolset-10 'gfortran -o output_file object_file..'
```

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

```
$ scl enable devtoolset-10 'bash'
```



NOTE

To verify the version of **gfortran** you are using at any point:

```
$ 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
```

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

```

Compile this source code on the command line by using the **gfortran** compiler from Red Hat Developer Toolset:

```
$ scl enable devtoolset-10 '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 run it:

```
$ ./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:

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

2.4. SPECIFICS OF GCC IN RED HAT DEVELOPER TOOLSET

Static linking of libraries

Certain more recent 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 creates an additional minor security risk as standard Red Hat Enterprise Linux errata do not change this code. If the need arises for developers to rebuild their applications due to this risk, Red Hat will communicate this using a security erratum.



IMPORTANT

Because of this additional security risk, developers are strongly advised not to statically link their entire application for the same reasons.

Specify libraries after object files when linking

In Red Hat Developer Toolset, libraries are linked using linker scripts which might specify some symbols through static archives. This is required to ensure compatibility with multiple versions of Red Hat Enterprise Linux. However, the linker scripts use the names of the respective shared object files. As a consequence, the linker uses different symbol handling rules than expected, and does not recognize symbols required by object files when the option adding the library is specified before options specifying the object files:

```
$ scl enable devtoolset-10 'gcc -lsomelib objfile.o'
```

Using a library from the Red Hat Developer Toolset in this manner results in the linker error message **undefined reference to symbol**. To prevent this problem, follow the standard linking practice, and specify the option adding the library after the options specifying the object files:

```
$ scl enable devtoolset-10 'gcc objfile.o -lsomelib'
```

Note that this recommendation also applies when using the base Red Hat Enterprise Linux version of **GCC**.

2.5. ADDITIONAL RESOURCES

For more information about the GNU Compiler Collections and its features, 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:

```
$ scl enable devtoolset-10 '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:

```
$ scl enable devtoolset-10 'man gfortran'
```

- *C++ Standard Library Documentation* – Documentation on the C++ standard library can be optionally installed:

```
# yum install devtoolset-10-libstdc++-docs
```

Once installed, HTML documentation is available at **/opt/rh/devtoolset-10/root/usr/share/doc/devtoolset-10-libstdC++-docs-10.2.1/html/index.html**.

Online Documentation

- [Red Hat Enterprise Linux 7 Developer Guide](#) – The *Developer Guide* for Red Hat Enterprise Linux 7 provides in-depth information about **GCC**.
- [Using the GNU Compiler Collection](#) – The upstream 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*](#) – An overview of Red Hat Developer Toolset and more information on how to install it on your system.
- [Chapter 4, *binutils*](#) – Instructions on using **binutils**, a collection of binary tools to inspect and manipulate object files and binaries.
- [Chapter 5, *elfutils*](#) – Instructions on using **elfutils**, a collection of binary tools to inspect and manipulate **ELF** files.
- [Chapter 6, *dwz*](#) – Instructions on using the **dwz** tool to optimize DWARF debugging information contained in **ELF** shared libraries and **ELF** executables for size.
- [Chapter 8, *GNU Debugger \(GDB\)*](#) – Instructions on debugging 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 4.2.1**. 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-10-make** package and is automatically installed with **devtoolset-10-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:

```
$ scl enable devtoolset-10 '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 **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 **make** as default:

```
$ scl enable devtoolset-10 'bash'
```

NOTE

To verify the version of **make** you are using at any point:

```
$ 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;
}
```

Build this source code using the implicit rules defined by the **make** utility from Red Hat Developer Toolset:

```
$ scl enable devtoolset-10 '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 the **-f** option:

```
$ make -f make_file
```

Describing the details of Makefile syntax is beyond the scope of this guide. See [GNU make](#), the upstream **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:

```
$ scl enable devtoolset-10 '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 $@

.o: .c
    $(CC) $(CFLAGS) $< -o $@
```

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

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

```
$ scl enable devtoolset-10 '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-10 'make clean'
rm -rf hello.o hello
```

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

3.4. ADDITIONAL RESOURCES

For more information about the **GNU make** tool and its features, 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:

```
$ scl enable devtoolset-10 '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:

```
$ scl enable devtoolset-10 'info make'
```

Online Documentation

- [GNU make](#) – The upstream **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](#) – An overview of Red Hat Developer Toolset and more information on how to install it on your system.
- [Chapter 2, GNU Compiler Collection \(GCC\)](#) – Instructions on using the **GNU Compiler Collection**, a portable compiler suite with support for a wide selection of programming languages.
- [Chapter 4, binutils](#) – Instructions on using **binutils**, a collection of binary tools to inspect and manipulate object files and binaries.

- [Chapter 5, *elfutils*](#) – Instructions on using **elfutils**, a collection of binary tools to inspect and manipulate **ELF** files.
- [Chapter 6, *dwz*](#) – Instructions on using the **dwz** tool to optimize DWARF debugging information contained in **ELF** shared libraries and **ELF** executables for size.
- [Chapter 8, *GNU Debugger \(GDB\)*](#) – Instructions on debugging 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 2.35**. 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.

Name	Description
strings	Displays printable character sequences in files.
strip	Discards all symbols from object files.

4.1. INSTALLING BINUTILS

In Red Hat Developer Toolset, **binutils** are provided by the **devtoolset-10-binutils** package and are automatically installed with **devtoolset-10-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:

```
$ scl enable devtoolset-10 'as option ... -o object_file source_file'
```

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

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

```
$ scl enable devtoolset-10 'bash'
```



NOTE

To verify the version of **as** you are using at any point:

```
$ which as
```

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

```
$ as -v
```

4.3. USING THE GNU LINKER

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

```
$ scl enable devtoolset-10 'ld option ... -o output_file object_file ...'
```

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

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

```
$ scl enable devtoolset-10 'bash'
```



NOTE

To verify the version of **ld** you are using at any point:

```
$ 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:

```
$ scl enable devtoolset-10 '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:

```
$ scl enable devtoolset-10 'objdump option ... object_file'
```

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

```
$ scl enable devtoolset-10 'bash'
```



NOTE

To verify the version of **binutils** you are using at any point:

```
$ 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. SPECIFICS OF BINUTILS IN RED HAT DEVELOPER TOOLSET

Static linking of libraries

Certain more recent 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 creates an additional minor security risk as standard Red Hat Enterprise Linux errata do not change this code. If the need arises for developers to rebuild their applications due to this risk, Red Hat will communicate this using a security erratum.



IMPORTANT

Because of this additional security risk, developers are strongly advised not to statically link their entire application for the same reasons.

Specify libraries after object files when linking

In Red Hat Developer Toolset, libraries are linked using linker scripts which might specify some symbols through static archives. This is required to ensure compatibility with multiple versions of Red Hat Enterprise Linux. However, the linker scripts use the names of the respective shared object files. As a consequence, the linker uses different symbol handling rules than expected, and does not recognize symbols required by object files when the option adding the library is specified before options specifying the object files:

```
$ scl enable devtoolset-10 'ld -lsomelib objfile.o'
```

Using a library from the Red Hat Developer Toolset in this manner results in the linker error message **undefined reference to symbol**. To prevent this problem, follow the standard linking practice, and specify the option adding the library after the options specifying the object files:

```
$ scl enable devtoolset-10 'ld objfile.o -lsomelib'
```

Note that this recommendation also applies when using the base Red Hat Enterprise Linux version of **binutils**.

4.6. ADDITIONAL RESOURCES

For more information about **binutils**, 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:

```
$ scl enable devtoolset-10 '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*](#) – An overview of Red Hat Developer Toolset and more information on how to install it on your system.
- [Chapter 5, *elfutils*](#) – Information on how to use **elfutils**, a collection of binary tools to inspect and manipulate ELF files.
- [Chapter 2, *GNU Compiler Collection \(GCC\)*](#) – 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 0.182**. This version is more recent than the version included in 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 (gABI)</i> and <i>processor-specific supplement ABI (psABI)</i> 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-10-elfutils** package and is automatically installed with **devtoolset-10-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 tool as follows:

```
$ scl enable devtoolset-10 '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:

```
$ scl enable devtoolset-10 'eu-objdump option ... object_file'
```

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

```
$ scl enable devtoolset-10 'bash'
```



NOTE

To verify the version of **elfutils** you are using at any point:

```
$ 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

For more information about **elfutils**, see the resources listed below.

See Also

- [Chapter 1, Red Hat Developer Toolset](#) – An overview of Red Hat Developer Toolset and more information on how to install it on your system.
- [Chapter 2, GNU Compiler Collection \(GCC\)](#) – Instructions on compiling programs written in C, C++, and Fortran.
- [Chapter 4, binutils](#) – Instructions on using **binutils**, a collection of binary tools to inspect and manipulate object files and binaries.
- [Chapter 6, dwz](#) – Instructions on using the **dwz** tool 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-10-dwz** package and is automatically installed with **devtoolset-10-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:

```
$ scl enable devtoolset-10 'dwz option... file_name'
```

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 **dwz** as default:

```
$ scl enable devtoolset-10 'bash'
```



NOTE

To verify the version of **dwz** you are using at any point:

```
$ 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

For more information about **dwz** and its features, 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:

```
$ scl enable devtoolset-10 'man dwz'
```

See Also

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

CHAPTER 7. ANNOBIN

The Annobin project consists of the **annobin** plugin and the **annockeck** program.

The **annobin** plugin scans the GNU Compiler Collection (GCC) command line, the compilation state, and the compilation process, and generates the ELF notes. The ELF notes record how the binary was built and provide information for the **annockeck** program to perform security hardening checks.

The security hardening checker is part of the **annockeck** program and is enabled by default. It checks the binary files to determine whether the program was built with necessary security hardening options and compiled correctly. **annockeck** is able to recursively scan directories, archives, and RPM packages for ELF object files.



NOTE

The files must be in ELF format. **annockeck** does not handle any other binary file types.

7.1. INSTALLING ANNOBIN

In Red Hat Developer Toolset, the **annobin** plugin and the **annockeck** program are provided by the **devtoolset-10-gcc** package and are installed as described in [Section 1.5.3, “Installing Optional Packages”](#).


7.2. USING ANNOBIN PLUGIN

To pass options to the **annobin** plugin with **gcc**, use:

```
$ scl enable devtoolset-10 'gcc -fplugin=annobin -fplugin-arg-annobin-option file-name'
```

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

```
$ scl enable devtoolset-10 'bash'
```



NOTE

To verify the version of **annobin** you are using at any point:

```
$ which annobin
```

Red Hat Developer Toolset’s **annobin** executable path will begin with **/opt**.

7.3. USING ANNOCHECK

To scan files, directories or RPM packages with the **annockeck** program:

```
$ scl enable devtoolset-10 'annockeck file-name'
```

**NOTE**

annocheck only looks for the ELF files. Other file types are ignored.

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

```
$ scl enable devtoolset-10 'bash'
```

**NOTE**

To verify the version of **annocheck** you are using at any point:

```
$ which annocheck
```

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

```
$ annocheck --version
```

7.4. ADDITIONAL RESOURCES

For more information about **annocheck**, **annobin** and its features, see the resources listed below.

Installed Documentation

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

```
$ scl enable devtoolset-10 'man annocheck'
```

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

```
$ scl enable devtoolset-10 'man annobin'
```

PART III. DEBUGGING TOOLS

CHAPTER 8. 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 9.2**. 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.

8.1. INSTALLING THE GNU DEBUGGER

In Red Hat Developer Toolset, the **GNU Debugger** is provided by the **devtoolset-10-gdb** package and is automatically installed with **devtoolset-10-toolchain** as described in [Section 1.5, “Installing Red Hat Developer Toolset”](#).

8.2. 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:

```
$ scl enable devtoolset-10 'gcc -g -o output_file input_file...'
```

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

```
$ scl enable devtoolset-10 'g++ -g -o output_file input_file...'
```

Example 8.1. 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;
}
```

Compile this program on the command line using **GCC** from Red Hat Developer Toolset with debugging information for the **GNU Debugger**:


```
$ scl enable devtoolset-10 '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:

```
# 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 8.2. Installing Debugging Information for the glibc Package

Install debugging information for the **glibc** package:

```
# 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
...
```

8.3. RUNNING THE GNU DEBUGGER

To run the **GNU Debugger** on a program you want to debug:

```
$ scl enable devtoolset-10 '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:

```
(gdb) 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-10 'bash'
```

**NOTE**

To verify the version of **gdb** you are using at any point:

```
$ 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 8.3. Running the gdb Utility on the fibonacci Binary File

This example assumes that you have successfully compiled the **fibonacci** binary file as shown in [Example 8.1, "Compiling a C Program With Debugging Information"](#).

Start debugging **fibonacci** with **gdb**:

```
$ scl enable devtoolset-10 'gdb fibonacci'
GNU gdb (GDB) Red Hat Enterprise Linux 8.2-2.el7
Copyright (C) 2017 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".
Type "show configuration" for configuration details.
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"...
Reading symbols from fibonacci...done.
(gdb)
```

8.4. LISTING SOURCE CODE

To view the source code of the program you are debugging:

```
(gdb) 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:

```
(gdb) list file_name:line_number
```

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

```
(gdb) list file_name:function_name
```

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

```
(gdb) set listsize number
```

Example 8.4. Listing the Source Code of the fibonacci Binary File

The **fibonacci.c** file listed in [Example 8.1, “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 }
```

8.5. SETTING BREAKPOINTS

Setting a New Breakpoint

To set a new breakpoint at a certain line:

```
(gdb) break file_name:line_number
```

You can also set a breakpoint on a certain function:

```
(gdb) break file_name:function_name
```

Example 8.5. Setting a New Breakpoint

This example assumes that you have compiled the **fibonacci.c** file listed in [Example 8.1, “Compiling a C Program With Debugging Information”](#) with debugging information.

Set a new breakpoint at line 10:

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

Listing Breakpoints

To display a list of currently set breakpoints:

```
(gdb) info breakpoints
```

Example 8.6. Listing Breakpoints

This example assumes that you have followed the instructions in [Example 8.5, “Setting a New Breakpoint”](#).

Display the list of currently set breakpoints:

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

Deleting Existing Breakpoints

To delete a breakpoint that is set at a certain line:

```
(gdb) clear line_number
```

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

```
(gdb) clear function_name
```

Example 8.7. Deleting an Existing Breakpoint

This example assumes that you have compiled the **fibonacci.c** file listed in [Example 8.1, “Compiling a C Program With Debugging Information”](#) with debugging information.

Set a new breakpoint at line 7:

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

Remove this breakpoint:

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

8.6. STARTING EXECUTION

To start an execution of the program you are debugging:

```
(gdb) run
```

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

```
(gdb) run argument...
```

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

Example 8.8. Executing the fibonacci Binary File

This example assumes that you have followed the instructions in [Example 8.5, "Setting a New Breakpoint"](#).

Execute the **fibonacci** binary file:

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

Breakpoint 1, main (argc=1, argv=0x7ffffffe4d8) at fibonacci.c:10
10      printf("%ld ", b);
```

8.7. 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:

```
(gdb) print variable_name
```

Example 8.9. Displaying the Current Values of Variables

This example assumes that you have followed the instructions in [Example 8.8, "Executing the fibonacci Binary File"](#) and the execution of the **fibonacci** binary stopped after reaching the breakpoint at line 10.

Display the current values of variables **a** and **b**:

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

8.8. CONTINUING EXECUTION

To resume the execution of the program you are debugging after it reached a breakpoint:

```
(gdb) continue
```

The execution stops again when another breakpoint is reached. To skip a certain number of breakpoints (typically when you are debugging a loop):

```
(gdb) continue number
```

The **gdb** utility also allows you to stop the execution after executing a single line of code:

```
(gdb) step
```

Finally, you can execute a certain number of lines:

```
(gdb) step number
```

Example 8.10. Continuing the Execution of the fibonacci Binary File

This example assumes that you have followed the instructions in [Example 8.8, "Executing the fibonacci Binary File"](#), and the execution of the **fibonacci** binary stopped after reaching the breakpoint at line 10.

Resume the execution:

```
(gdb) continue  
Continuing.
```

```
Breakpoint 1, main (argc=1, argv=0x7ffffffe4d8) at fibonacci.c:10  
10     printf("%ld ", b);
```

The execution stops the next time the breakpoint is reached.

Execute the next three lines of code:

```
(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
```

8.9. ADDITIONAL RESOURCES

For more information about the **GNU Debugger** and all its features, see the resources listed below.

Installed Documentation

Installing the **devtoolset-10-gdb-doc** package provides the following documentation in HTML and PDF formats in the `/opt/rh/devtoolset-10/root/usr/share/doc/devtoolset-10-gdb-doc-9.2` directory:

- The *Debugging with GDB* book, which is a copy of the upstream material with the same name. The version of this document exactly corresponds to the version of **GDB** available in Red Hat Developer Toolset.
- The *GDB's Obsolete Annotations* document, which lists the obsolete **GDB** level 2 annotations.

Online Documentation

- [Red Hat Enterprise Linux 7 Developer Guide](#) – The *Developer Guide* for Red Hat Enterprise Linux 7 provides more information on the **GNU Debugger** and debugging.
- [GDB Documentation](#) – The upstream **GDB** documentation includes the *GDB User Manual* and other reference material.

See Also

- [Chapter 1, Red Hat Developer Toolset](#) – An overview of Red Hat Developer Toolset and more information on how to install it on your system.
- [Chapter 2, GNU Compiler Collection \(GCC\)](#) – Further information on how to compile programs written in C, C++, and Fortran.
- [Chapter 9, strace](#) – Instructions on using the **strace** utility to monitor system calls that a program uses and signals it receives.
- [Chapter 11, memstomp](#) – Instructions on using the **memstomp** utility to identify calls to library functions with overlapping memory regions that are not allowed by various standards.

CHAPTER 9. 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 5.7**.

9.1. INSTALLING STRACE

In Red Hat Enterprise Linux, the **strace** utility is provided by the **devtoolset-10-strace** package and is automatically installed with **devtoolset-10-toolchain** as described in [Section 1.5, “Installing Red Hat Developer Toolset”](#).

9.2. USING STRACE

To run the **strace** utility on a program you want to analyze:

```
$ scl enable devtoolset-10 '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:

```
$ scl enable devtoolset-10 'strace -p process_id'
```

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 **strace** as default:

```
$ scl enable devtoolset-10 'bash'
```

NOTE

To verify the version of **strace** you are using at any point:

```
$ 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
```

9.2.1. 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-10 'strace -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 8.1, “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. Run the **strace** utility on this file and redirect the trace output to **fibonacci.log**:

```
$ scl enable devtoolset-10 '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.

9.2.2. 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-10 '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 9.1, “Commonly Used Values of the -e Option”](#). For a detailed description of all available values, see the **strace(1)** manual page.

Table 9.1. Commonly Used Values of the -e Option

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

Note that the syntax **-e *expression*** is a shorthand for the full form **-e trace=*expression***.

Example 9.2. Tracing Selected System Calls

Consider the **employee** file from [Example 11.1, “Using memstomp”](#). Run the **strace** utility on this executable file and trace only the **mmap** and **munmap** system calls:

```
$ scl enable devtoolset-10 '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 +++

```

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 **strace** utility with the **-t** command line option:

```
$ scli enable devtoolset-10 'strace -t program argument...'
```

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

```
$ scli enable devtoolset-10 '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:

```
$ scli enable devtoolset-10 'strace -r program argument...'
```

Example 9.3. Displaying Time Stamps

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

```

$ scli enable devtoolset-10 '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 +++
```

9.2.4. 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-10 'strace -c program argument...'
```

Example 9.4. Displaying a Summary

Consider an executable file named **lsblk**. Run the **strace** utility on this file and display a trace summary:

```
$ scl enable devtoolset-10 '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     0 munmap
  0.00   0.000000     0    148     0 read
  0.00   0.000000     0     1     0 write
  0.00   0.000000     0   258     0 close
  0.00   0.000000     0    37     0 stat
...
-----
100.00   0.000068           1790    35 total
```

9.2.5. Tampering with System Call Results

Simulating errors returned from system calls can help identify missing error handling in programs.

To make a program receive a generic error as the result of a particular system call, run the **strace** utility with the **-e fault=** option and supply the system call:

```
$ scl enable devtoolset-10 'strace -e fault=syscall program argument...'
```

To specify the error type or return value, use the **-e inject=** option:

```
$ scl enable devtoolset-10 'strace -e inject=syscall:error=error-type program argument'
$ scl enable devtoolset-10 'strace -e inject=syscall:retval=return-value program argument'
```

Note that specifying the error type and return value is mutually exclusive.

Example 9.5. Tampering with System Call Results

Consider an executable file named **lsblk**. Run the **strace** utility on this file and make the **mmap()** system call return an error:

```
-
```

```

$ scl enable devtoolset-10 'strace -e fault=mmap:error=EPERM lsblk > /dev/null'
execve("/usr/bin/lsblk", ["lsblk"], 0x7fff1c0e02a0 /* 54 vars */) = 0
brk(NULL)                               = 0x55d9e8b43000
mmap(NULL, 8192, PROT_READ|PROT_WRITE, MAP_PRIVATE|MAP_ANONYMOUS, -1, 0) = -
1 EPERM (Operation not permitted) (INJECTED)
writev(2, [{iov_base="lsblk", iov_len=5}, {iov_base=": ", iov_len=2}, {iov_base="error while loading
shared libra...", iov_len=36}, {iov_base=": ", iov_len=2}, {iov_base="", iov_len=0}, {iov_base="",
iov_len=0}, {iov_base="cannot create cache for search p...", iov_len=35}, {iov_base=": ",
iov_len=2}, {iov_base="Cannot allocate memory", iov_len=22}, {iov_base="\n", iov_len=1}],
10lsblk: error while loading shared libraries: cannot create cache for search path: Cannot allocate
memory
) = 105
exit_group(127)                          = ?
+++ exited with 127 +++

```

9.3. ADDITIONAL RESOURCES

For more information about **strace** and its features, 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:

```
$ scl enable devtoolset-10 'man strace'
```

See Also

- [Chapter 1, Red Hat Developer Toolset](#) – An overview of Red Hat Developer Toolset and more information on how to install it on your system.
- [Chapter 10, ltrace](#) – Instructions on tracing program library calls using the **ltrace** tool.
- [Chapter 8, GNU Debugger \(GDB\)](#) – Instructions on debugging programs written in C, C++, and Fortran.
- [Chapter 11, memstomp](#) – Instructions on using the **memstomp** utility to identify calls to library functions with overlapping memory regions that are not allowed by various standards.

CHAPTER 10. 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.

10.1. INSTALLING LTRACE

In Red Hat Enterprise Linux, the **ltrace** utility is provided by the **devtoolset-10-ltrace** package and is automatically installed with **devtoolset-10-toolchain** as described in [Section 1.5, “Installing Red Hat Developer Toolset”](#).

10.2. USING LTRACE

To run the **ltrace** utility on a program you want to analyze:

```
$ scl enable devtoolset-10 '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:

```
$ scl enable devtoolset-10 'ltrace -p process_id'
```

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 **ltrace** as default:

```
$ scl enable devtoolset-10 'bash'
```



NOTE

To verify the version of **ltrace** you are using at any point:

```
$ 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
```

10.2.1. 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-10 'ltrace -o file_name program argument...'
```

Replace *file_name* with the name of the file.

Example 10.1. Redirecting Output to a File

Consider a slightly modified version of the **fibonacci** file from [Example 8.1, “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. Run the **ltrace** utility on this file and redirect the trace output to **fibonacci.log**:

```
$ scl enable devtoolset-10 '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.

10.2.2. 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-10 '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:

```
$ scl enable devtoolset-10 'ltrace -e malloc+free-@libc.so* program'
```

Example 10.2. Tracing Selected Library Calls

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

```
$ scl enable devtoolset-10 '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.

10.2.3. 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-10 'ltrace -t program argument...'
```

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

```
$ scl enable devtoolset-10 '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-10 'ltrace -r program argument...'
```

Example 10.3. Displaying Time Stamps

Consider the **pwd** command. Run the **ltrace** utility on this program and include time stamps in the output:

```
$ scl enable devtoolset-10 '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)
```

10.2.4. 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-10 'ltrace -c program argument...'
```

Example 10.4. Displaying a Summary

Consider the **lsblk** command. Run the **ltrace** utility on this program and display a trace summary:

```
$ scl enable devtoolset-10 '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
```

10.3. ADDITIONAL RESOURCES

For more information about **ltrace** and its features, 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:

```
$ scl enable devtoolset-10 '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](#) – An overview of Red Hat Developer Toolset and more information on how to install it on your system.
- [Chapter 9, strace](#) – Instructions on tracing program system calls using the **strace** tool.
- [Chapter 8, GNU Debugger \(GDB\)](#) – Instructions on debugging programs written in C, C++, and Fortran.
- [Chapter 11, memstomp](#) – Instructions on using the **memstomp** utility to identify calls to library functions with overlapping memory regions that are not allowed by various standards.

CHAPTER 11. 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 11.1, “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 11.1. 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
wmempcpy	The wide-character equivalent of the memcpy() function that copies n 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 n 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 n 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.

11.1. INSTALLING MEMSTOMP

In Red Hat Developer Toolset, the **memstomp** utility is provided by the **devtoolset-10-memstomp** package and is automatically installed with **devtoolset-10-toolchain** as described in [Section 1.5, “Installing Red Hat Developer Toolset”](#).

11.2. USING MEMSTOMP

To run the **memstomp** utility on a program you want to analyze:

```
$ scl enable devtoolset-10 '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-10 '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-10 '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 8.2, “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-10 'bash'
```

Example 11.1. 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**:

```
$ scl enable devtoolset-10 'gcc -rdynamic -g -o employee employee.c'
```

To identify erroneous function calls with overlapping memory regions:

```
$ scl enable devtoolset-10 '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,
```

11.3. ADDITIONAL RESOURCES

For more information about **memstomp** and its features, 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:

```
$ scl enable devtoolset-10 'man memstomp'
```

See Also

- [Chapter 1, Red Hat Developer Toolset](#) – An overview of Red Hat Developer Toolset and more information on how to install it on your system.
- [Chapter 8, GNU Debugger \(GDB\)](#) – Instructions on debugging programs written in C, C++, and Fortran.
- [Chapter 9, strace](#) – Instructions on using the `strace` utility to monitor system calls that a program uses and signals it receives.
- [Chapter 13, Valgrind](#) – Instructions on using the **Valgrind** tool 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 12. 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 4.4**. 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 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.

12.1. INSTALLING SYSTEMTAP

In Red Hat Developer Toolset, **SystemTap** is provided by the **devtoolset-10-systemtap** package and is automatically installed with **devtoolset-10-perftools** as described in [Section 1.5, “Installing Red Hat Developer Toolset”](#).

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

```
$ scl enable devtoolset-10 '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 7 SystemTap Beginners Guide](#).

12.2. USING SYSTEMTAP

To execute any of the tools that are part of **SystemTap**:

```
$ scli enable devtoolset-10 'tool option...'
```

See [Table 12.1, “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:

```
$ scli enable devtoolset-10 'stap option... argument...'
```

Note that you can execute any command using the **scli** 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:

```
$ scli enable devtoolset-10 'bash'
```



NOTE

To verify the version of **SystemTap** you are using at any point:

```
$ 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
```

12.3. ADDITIONAL RESOURCES

For more information about **SystemTap** and its features, 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:

```
$ scli enable devtoolset-10 '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:

```
$ scli enable devtoolset-10 'man staprun'
```

Online Documentation

- [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 7 SystemTap Tapset Reference](#) – The *SystemTap Tapset Reference* for Red Hat Enterprise Linux 7 provides further details about **SystemTap**.
- [The SystemTap Documentation](#) – The **SystemTap** documentation provides further documentation about **SystemTap**, and numerous examples of the **SystemTap** scripts.

See Also

- [Chapter 1, Red Hat Developer Toolset](#) – An overview of Red Hat Developer Toolset and more information on how to install it on your system.
- [Chapter 13, Valgrind](#) – Instructions on using the **Valgrind** tool 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, OProfile](#) – Instructions on using the **OProfile** tool to determine which sections of code consume the greatest amount of **CPU** time and why.
- [Chapter 15, Dyninst](#) – Instructions on using the **Dyninst** library to instrument a user-space executable.

CHAPTER 13. 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 13.1, “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 3.16.1**. 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 13.1. 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.

13.1. INSTALLING VALGRIND

In Red Hat Developer Toolset, **Valgrind** is provided by the **devtoolset-10-valgrind** package and is automatically installed with **devtoolset-10-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.

13.2. USING VALGRIND

To run any of the **Valgrind** tools on a program you want to profile:

```
$ scl enable devtoolset-10 'valgrind --tool=tool program argument...'
```

See [Table 13.1, “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:

```
$ scl enable devtoolset-10 'valgrind --tool=cachegrind program argument...'
```

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

```
$ scl enable devtoolset-10 'bash'
```



NOTE

To verify the version of Valgrind you are using at any point:

```
$ 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
```

13.3. ADDITIONAL RESOURCES

For more information about **Valgrind** and its features, 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:

```
$ scl enable devtoolset-10 'man valgrind'
```

- *Valgrind Documentation* – **HTML** documentation for **Valgrind** is located at **/opt/rh/devtoolset-10/root/usr/share/doc/devtoolset-10-valgrind-3.16.1/html/index.html**.

Online Documentation

- [Red Hat Enterprise Linux 7 Developer Guide](#) – The *Developer Guide* for Red Hat Enterprise Linux 7 provides more information about **Valgrind** and its **Eclipse** plug-in.
- [Red Hat Enterprise Linux 7 Performance Tuning Guide](#) – The *Performance Tuning Guide* for Red Hat Enterprise Linux 7 provide more detailed information about using **Valgrind** to profile applications.

See Also

- [Chapter 1, *Red Hat Developer Toolset*](#) – An overview of Red Hat Developer Toolset and more information on how to install it on your system.
- [Chapter 11, *memstomp*](#) – Instructions on using the **memstomp** utility to identify calls to library functions with overlapping memory regions that are not allowed by various standards.
- [Chapter 12, *SystemTap*](#) – An introduction to the **SystemTap** tool and instructions on how to use it to monitor the activities of a running system.
- [Chapter 14, *OProfile*](#) – Instructions on using the **OProfile** tool to determine which sections of code consume the greatest amount of CPU time and why.
- [Chapter 15, *Dyninst*](#) – Instructions on using the **Dyninst** library to instrument a user-space executable.

CHAPTER 14. 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 14.1, “Tools Distributed with OProfile for Red Hat Developer Toolset”](#).

OProfile profiles an application without adding any instrumentation by recording the details of every *n*th 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 1.4.0**.

Table 14.1. Tools Distributed with OProfile for Red Hat Developer Toolset

Name	Description
opperf	Records samples either for a single process or system-wide using the Linux Performance Events subsystem.
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.

14.1. INSTALLING OPROFILE

In Red Hat Developer Toolset, **OProfile** is provided by the **devtoolset-10-oprofile** package and is automatically installed with **devtoolset-10-perftools** as described in [Section 1.5, “Installing Red Hat Developer Toolset”](#).

14.2. USING OPROFILE

To run any of the tools that are distributed with **OProfile**:

```
# scl enable devtoolset-10 'tool option...'
```

See [Table 14.1, “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:

```
$ scl enable devtoolset-10 'ophelp -X'
```

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 **OProfile** as default:

```
$ scl enable devtoolset-10 'bash'
```



NOTE

To verify the version of **OProfile** you are using at any point:

```
$ 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
```

14.3. ADDITIONAL RESOURCES

For more information about **OProfile** and its features, 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:

```
$ scl enable devtoolset-10 '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:

```
scl enable devtoolset-10 'man tool'
```

Online Documentation

- [Red Hat Enterprise Linux 7 Developer Guide](#) – The *Developer Guide* for Red Hat Enterprise Linux 7 provides more information on **OProfile**.
- [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*](#) – An overview of Red Hat Developer Toolset and more information on how to install it on your system.
- [Chapter 12, *SystemTap*](#) – An introduction to **SystemTap** and instructions on how to use it to monitor the activities of a running system.
- [Chapter 13, *Valgrind*](#) – Instructions on using the **Valgrind** tool 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 15, *Dyninst*](#) – Instructions on using the **Dyninst** library to instrument a user-space executable.

CHAPTER 15. 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 10.2.1**.

15.1. INSTALLING DYNINST

In Red Hat Developer Toolset, the **Dyninst** library is provided by the **devtoolset-10-dyninst** package and is automatically installed with **devtoolset-10-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-10-toolchain** package.

If you intend to write a custom instrumentation for binaries, install the relevant header files:

```
# yum install devtoolset-10-dyninst-devel
```

You can also install **API** documentation for this library:

```
# yum install devtoolset-10-dyninst-doc
```

For a complete list of documents that are included in the **devtoolset-10-dyninst-doc** package, see [Section 15.3, “Additional Resources”](#). For detailed instructions on how to install optional packages to your system, see [Section 1.5, “Installing Red Hat Developer Toolset”](#).

15.2. USING DYNINST

15.2.1. 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:

```
$ scl enable devtoolset-10 "stap --dyninst -c 'command' option... argument..."
```

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

```
$ scl enable devtoolset-10 "stap --dyninst -x process_id option... argument..."
```

See [Chapter 12, *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 7.

Example 15.1. 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. Compile this program on the command line using the **g++** compiler from Red Hat Developer Toolset:

```
$ scl enable devtoolset-10 '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. Run this script on the **exercise** binary file:

```
$ scl enable devtoolset-10 "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.
```

15.2.2. Using Dyninst as a Stand-alone Library

Before using the **Dyninst** library as a part of your application, set the value of the **DYNINSTAPI_RT_LIB** environment variable to the path to the runtime library file:


```
$ export DYNINSTAPI_RT_LIB=/opt/rh/devtoolset-10/root/usr/lib64/dyninst/libdyninstAPI_RT.so
```

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

[Example 15.2, “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 15.3, “Additional Resources”](#).

Example 15.2. Using Dyninst as a Stand-alone Application

Consider the **exercise.C** source file from [Example 15.1, “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 *applImage) {
    int zero = 0;
    counter = app->malloc(*applImage->findType("int"));
    counter->writeValue(&zero);
}

bool interceptfunc(BPatch_process *app,
                  BPatch_image *applImage,
                  char *funcName) {
    BPatch_Vector<BPatch_function *> func;
    applImage->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 *applImage = app->getImage();
    createCounter(app, applImage);
    fprintf(stderr, "Finding function %s(): ", argv[2]);
    BPatch_Vector<BPatch_function*> countFuncs;
    fprintf(stderr, "OK\nInstrumenting function %s(): ", argv[2]);
    interceptfunc(app, applImage, 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-10/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:

```
$ scl enable devtoolset-10 make
g++ -g -I/opt/rh/devtoolset-10/root/usr/include/dyninst count.C -c
g++ -g -I/opt/rh/devtoolset-10/root/usr/include/dyninst count.o -L /opt/rh/devtoolset-10/root/usr/lib64/dyninst -ldyninstAPI -o count
g++ -g -I/opt/rh/devtoolset-10/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-10/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.

15.3. ADDITIONAL RESOURCES

For more information about Dyninst and its features, see the resources listed below.

Installed Documentation

The `devtoolset-10-dyninst-doc` package installs the following documents in the `/opt/rh/devtoolset-10/root/usr/share/doc/devtoolset-10-dyninst-doc-10.2.1/` 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 15.1, "Installing Dyninst"](#).

Online Documentation

- [Dyninst Home Page](#) – The project home page provides links to additional documentation and related publications.
- [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 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](#) – An overview of Red Hat Developer Toolset and more information on how to install it on your system.
- [Chapter 12, SystemTap](#) – An introduction to **SystemTap** and instructions on how to use it to monitor the activities of a running system.

- [Chapter 13, Valgrind](#) – Instructions on using the **Valgrind** tool 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, OProfile](#) – Instructions on using the **OProfile** tool to determine which sections of code consume the greatest amount of **CPU** time and why.

PART V. COMPILER TOOLSETS

CHAPTER 16. COMPILER TOOLSETS DOCUMENTATION

The descriptions of the three compiler toolsets:

- LLVM Toolset
- Go Toolset
- Rust Toolset

have been moved to a separate documentation set under [Red Hat Developer Tools](#).

PART VI. GETTING HELP

CHAPTER 17. 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 10.1 Release Notes](#) – The *Release Notes* for Red Hat Developer Toolset 10.1 contain more information.
- [Using Red Hat Software Collections Container Images](#) – The *Using Red Hat Software Collections Container Images* provides instructions for obtaining, configuring, and using container images that are shipped with Red Hat Software Collections, including the Red Hat Developer Toolset container images.
- [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 7 Developer Guide](#) – The *Developer Guide* for Red Hat Enterprise Linux 7 provides more information about libraries and runtime support, compiling and building, debugging, and profiling.
- [Red Hat Enterprise Linux 7 Installation Guide](#) – The *Installation Guide* for Red Hat Enterprise Linux 7 explains how to obtain, install, and update the system.
- [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 18. 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).


18.1. 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:

```
# yum install sos
```

To generate the report:

```
# 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>.

18.2. 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.

18.3. 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.

18.4. 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 10.0

The sections below document features and compatibility changes introduced Red Hat Developer Toolset 10.0.

A.1. CHANGES IN GCC

Red Hat Developer Toolset 10.0 is distributed with **GCC 10.2.1**.

The following features have been added or modified since the previous release of Red Hat Developer Toolset:

General improvements

- New built-in functions:
 - The **__has_builtin** built-in preprocessor operator can now be used to query support for built-in GCC functions.
 - The **__builtin_roundeven** function has been added for the corresponding function from ISO/IEC TS 18661.
- New command-line options:
 - **-fallocation-dce** removes unnecessary pairs of the **new** and **delete** operators.
 - **-fprofile-partial-training** informs the compiler that it should optimize only the size of code paths that are covered by the training run.
 - **-fprofile-reproducible** controls the level of reproducibility of profiles gathered by **-fprofile-generate**. Using this option you can rebuild a program with the same outcome. It is useful for distribution packages, for example.
 - Experimental: the new **-fanalyzer** option enables a new static analysis pass and associated warnings. This pass examines paths in the code to detect various common errors, for example, double-free bugs. Note that this option is available only for code written in C.
- Interprocedural optimization improvements:
 - The interprocedural scalar replacement of aggregates (IPA-SRA) pass was re-implemented to work at link time. It can now also remove computing and returning unused return values.
 - The **-finline-functions** option is now enabled at the optimization level **-O2**. This option now reduces code size better. Inliner heuristics now work faster to avoid negative impact on **-flto -O2** compile times.
 - Inliner heuristics and function cloning can now use value-range information to predict the effectiveness of individual transformations.
 - During link-time optimization the C++ One Definition Rule is now used to increase precision of type-based alias analysis.
- Link-time optimization improvements:
 - The parallel phase of the Link Time Optimization (LTO) can now automatically detect a running **jobserver** of the **make** tool.

- Profile-driven optimization improvements:
 - Profile maintenance during compilation and hot or cold code partitioning have been improved.
- When GCC prints a warning, the option that controls the warning is now displayed as a hyperlink which you can click and see the documentation for the particular warning. To control this behavior, use the **-fdiagnostics-urls** option.

Language-specific improvements

- Several newly implemented OpenMP 5.0 features have been added. For example, the **conditional lastprivate** clause, **scan** and **loop** directives, **order(concurrent)** and **use_device_addr** clauses support, **if** clauses on **simd** constructs or partial support for the **declare variant** directive.

Notable changes related to languages include:

C family

- New attributes:
 - The **access** function and type attribute has been added. It describes how a function accesses objects passed to it by a pointer or reference and associates such arguments with integer arguments that denote the objects' sizes. You can also use this attribute to enable the detection of invalid accesses by user-defined functions, such as those diagnosed by the **-Wstringop-overflow** option.
- New warnings:
 - **-Wstring-compare**, enabled by the **-Wextra** option, warns about equality and inequality expressions between zero and the result of a call to **strcmp** or **strncmp** when such expressions evaluate to a constant because the length of one argument is greater than the size of the array pointed to by the other argument.
 - **-Wzero-length-bounds**, enabled by the **-Warray-bounds** option, warns about accesses to elements of zero-length arrays that might overlap with other members of the same object.
- Enhancements to existing warnings:
 - **-Warray-bounds** now detects more out-of-bounds accesses to member arrays and elements of zero-length arrays.
 - **-Wformat-overflow** now fully uses string length information calculated by the **strlen** optimization pass.
 - **-Wrestrict** detects overlapping accesses to dynamically allocated objects.
 - **-Wreturn-local-addr** diagnoses more instances of return statements returning addresses of automatic variables.
 - **-Wstringop-overflow** detects more out-of-bounds stores to member arrays including zero-length arrays, dynamically allocated objects, and variable length arrays. It now also detects more instances of reads of unterminated character arrays by the built-in string functions. In addition, the warning now detects out-of-bounds accesses by calls to user-defined functions declared with the new attribute **access**.
 - **-Warith-conversion** re-enables warnings from **-Wconversion**, **-Wfloat-conversion**, and -

Wsign-conversion. These warnings are now disabled by default for expressions where the result of an arithmetic operation will not fit into the target type because of promotion, but where the operands of the expressions fit into the target type.

C

- Several new features from the upcoming C2X revision of the ISO C standard are now supported. To enable them, use **-std=c2x** and **-std=gnu2x**. Some of these features are also supported as extensions when you compile with earlier C versions. Some features were previously supported as extensions and now have been added to the C standard. They are enabled by default in C2X mode and not diagnosed with **-std=c2x -Wpedantic**.
 - The **[[]]** attribute syntax is now supported.
 - In C2X mode, empty parentheses in a function definition give that function a type with a prototype for subsequent calls. Other old-style function definitions are diagnosed by default in C2X mode.
- GCC now defaults to **-fno-common**. As a result, global variable accesses are more efficient on various targets. However, in the C language, global variables with multiple tentative definitions now cause linker errors. With the **-fcommon** option such definitions are silently merged during linking.

C++

- The following C++20 features have been implemented:
 - Concepts, including the P0734R0, P0857R0, P1084R2, P1141R2, P0848R3, P1616R1, and P1452R2 proposals
 - P1668R1: Permitting unevaluated inline-assembly in **constexpr** functions
 - P1161R3: Deprecate **a[b,c]**
 - P0848R3: Conditionally trivial special member functions
 - P1091R3: Structured binding extensions
 - P1143R2: Adding the **constexpr** keyword
 - P1152R4: Deprecating **volatile**
 - P0388R4: Permit conversions to arrays of unknown bound
 - P0784R7: More **constexpr** containers (**constexpr new**)
 - P1301R4: **[[nodiscard("with reason")]]**
 - P1814R0: Class template argument deduction for alias templates
 - P1816R0: Class template argument deduction for aggregates
 - P0960R3: Parenthesized initialization of aggregates
 - P1331R2: Permitting trivial default initialization in **constexpr** contexts
 - P1327R1: Allowing **dynamic_cast** and polymorphic **typeid** in constant expressions

- P0912R5: Coroutines. The **-fcoroutines** option enables the support for coroutines. To learn more about the above-mentioned proposals, see [C++ Standards Support in GCC](#).
- Several C++ defect reports have been resolved. You can find the overall defect report status on the [C++ Defect Report Support in GCC](#) page.
- New warnings:
 - G++ can now detect modifying constant objects in the **constexpr** evaluation, which is undefined behavior.
 - Memory consumption of the compiler, when performing the **constexpr** evaluation, has been reduced.
 - The **noexcept** specifier is now properly treated as a *complete-class context*.
 - You can now use the **deprecated** attribute on namespaces.

Runtime library **libstdc++**

- The following experimental C++2a support features have been improved:
 - Library concepts in **<concepts>** and **<iterator>**
 - Constrained algorithms in **<ranges>**, **<algorithm>**, and **<memory>**
 - New algorithms **shift_left** and **shift_right**
 - The **std::span** class template
 - Three-way comparisons in **<compare>** and throughout the library
 - **Constexpr** support in **<algorithm>** and other places
 - **<stop_token>** and **std::jthread**
 - **std::atomic_ref** and **std::atomic<floating point>**
 - Integer comparison functions (for example, **cmp_equal**, **cmp_less**, and other functions)
 - **std::ssize** and **std::to_array**
 - **std::construct_at**, **std::destroy**, and **constexpr std::allocator**
 - Mathematical constants in **<numbers>**
- The random number generator **std::random_device** now supports RDSEED.

Fortran

- The compiler now rejects mismatches between actual and dummy argument lists in a single file and prints an error. To turn these errors into warnings, use the new **-fallow-argument-mismatch** option. This option is implied with **std=legacy**. The **-Wargument-mismatch** option has been removed.
- The binary, octal, and hexadecimal (BOZ) literal constants have been improved and now conform better to the Fortran 2008 and 2018 standards. In these Fortran standards, BOZ literal constants do not have types or kinds. With this enhancement, documented and undocumented

extensions to the Fortran standards now emit errors during compilation. You can enable some of these extensions with the **-fallow-invalid-boz** option, which turns errors into warnings and the code is compiled as with the earlier GFortran.

Target-specific improvements

Changes to architecture and processor support include:

AMD64 and Intel 64

- GCC now supports the Cooper Lake Intel CPU. To enable it, use the **-march=cooperlake** option, which enables the AVX512BF16 ISA extensions.
- GCC now supports the Tiger Lake Intel CPU. To enable it, use the **-march=tigerlake** option, which enables the MOVDIRI MOVDIR64B AVX512VP2INTERSECT ISA extensions.

A.2. CHANGES IN BINUTILS

Red Hat Developer Toolset 10.0 is distributed with **binutils 2.35**.

The following features have been added or modified since the previous release of Red Hat Developer Toolset:

The assembler

- The **.symver** directive has been extended to update the visibility of the original symbol and assign one original symbol to different versioned symbols.
- The Intel SERIALIZE and TSXLDTRK instructions are now supported.
- The **-mfence-after-load=**, **-mfence-before-indirect-branch=**, and **-mfence-before-ret=** options have been added to the x86 assembler to help mitigate [CVE-2020-0551](#).
- The **--gdwarf-5** option has been added to the assembler to generate DWARF 5 debug output if such output is generated. Also, it is now possible to generate version 5 **.debug_line** sections.
- The **-malign-branch-boundary=NUM**, **-malign-branch=TYPE[+TYPE...]**, **-malign-branch-prefix-size=NUM**, and **-mbranches-within-32B-boundaries** options have been added to the x86 assembler to align branches within a fixed boundary with segment prefixes or the NOP instructions.
- The **--gdwarf-cie-version** command-line flag has been added. This flag controls which version of DWARF CIE (Common Information Entries) the assembler creates.

The linker

- The command-line options **--export-dynamic-symbol** and **--export-dynamic-symbol-list** have been added to make symbols dynamic.
- The **-Map=filename** command-line option has been extended. If **filename** is a directory, the linker will create the **filename/output-filename.map** file.
- The **--warn-textrel** command-line option has been added to warn about **DT_TEXTREL** being set in a position-independent executable or shared object.
- The command-line options **--enable-non-contiguous-regions** and **--enable-non-contiguous-regions-warnings** have been added.

- Relative path names in **INPUT()** and **GROUP()** directives in linker scripts are now searched in relation to the directory of the linker script before other search paths.
- The **-z start-stop-visibility=...** command-line option has been added to control the visibility of synthetic **__start_SECNAME** and **__stop_SECNAME** symbols.
- The **--dependency-file** command-line option has been added to write a Make-style dependency file listing the input files consulted by the linker, like the files written by the compiler **-M** and **-MP** options.
- The **ld** check for the **PHDR segment not covered by LOAD segment** error is more effective now. The check now catches cases that earlier versions of **ld** incorrectly allowed. If you see this error, it is likely you are linking with a bad linker script or the binary you are building is not intended to be loaded by a dynamic loader. In the latter case, the **--no-dynamic-linker** option is appropriate.
- The **--no-print-map-discarded** command-line option has been added.

Other binary utilities

- The **readelf** tool now displays symbol names differently when wide mode is not enabled. If the name is too long, it will be shortened and the last five characters replaced with "[...]". You can restore the old behaviour with the **-T** or **--silent-truncation** option.
- The **readelf** tool now has the **-L** or **--lint** or **--enable-checks** option, which enables warning messages about possible problems with the file(s) being examined. For example, with this option enabled, **readelf** will check for zero-sized sections, which are allowed by the ELF standard but might be potentially dangerous if the user was expecting them to actually contain something.
- **Binutils** now supports **debuginfod**, an HTTP server for distributing ELF/DWARF debugging information as well as source code. When built with **debuginfod**, **readelf** and **objdump** can automatically query **debuginfod** servers for separate debugging files when such files otherwise cannot be found. To build **binutils** with **debuginfod**, pass the **--with-debuginfod** configure option. This requires **libdebuginfod**, the **debuginfod** client library. **debuginfod** is distributed with **elfutils**, starting with version 0.178. For more information, see <https://sourceware.org/elfutils>.
- The **--output** option has been added to the **ar** program. This option can be used to specify the output directory when extracting members from an archive.
- The **--keep-section** option has been added to **objcopy** and **strip**. This option keeps the specified section from being removed.
- The **--source-comment[=<txt>]** option has been added to **objdump**. It provides a prefix to source code lines displayed in a disassembly.
- The **--set-section-alignment <section-name>=<align>** option has been added to **objcopy** to allow the changing of section alignments.
- The **--verilog-data-width** option has been added to **objcopy** for Verilog targets to control the width of data elements in the Verilog hexadecimal format.
- The separate debug information file options of **readelf** (**--debug-dump=links** and **--debug-dump=follow**) and **objdump** (**--dwarf=links** and **--dwarf=follow-links**) will now display or follow (or both) multiple links if a file has more than one such link. This usually happens when the **GCC -gsplit-dwarf** option is used.

In addition, the **--dwarf=follow-links** option for **objdump** now also affects its other display options. For example, when combined with the **--syms** option, it will cause the symbol tables in any linked debug information files to also be displayed. When combined with the **--disassemble** option, the **--dwarf= follow-links** option will ensure that any symbol tables in the linked files are read and used when disassembling code in the main file.

- Dumping types encoded in the Compact Type Format are now supported in **objdump** and **readelf**.

A.3. CHANGES IN ELFUTILS

Red Hat Developer Toolset 10.0 is distributed with **elfutils 0.180**.

The following features have been added or modified since the previous release of Red Hat Developer Toolset:

- The new **eu-elfclassify** tool has been added to analyze ELF objects.
- The new **debuginfod** server with a client tool and library has been added. **debuginfod** indexes and automatically fetches ELF, DWARF, and source from files and RPM archives through HTTP.
- **libebl** is now directly compiled into **libdw.so**.
- **eu-readelf** now has multiple new flags for notes, section numbering, and symbol tables.
- **libdw** has improved multithreading support.
- **libdw** now supports additional GNU DWARF extensions.
- Better support for debug info for code built with GCC LTO (link time optimization). The **eu-readelf** and **libdw** utilities can now read and handle **.gnu.debuglto_** sections, and correctly resolve file names for functions that are defined across compile units (CUs).
- The **eu-nm** utility now explicitly identifies weak objects as **V** and common symbols as **C**.

A.4. CHANGES IN GDB

Red Hat Developer Toolset 10.0 is distributed with **GDB 9.2**.

The following features have been added or modified since the previous release of Red Hat Developer Toolset:

New features

- The **debuginfod** server is now supported. This is an HTTP server for distributing ELF and DWARF debugging information and source code.

New convenience variables and functions

- **\$_gdb_major**
\$_gdb_minor

With these new convenience variables for testing the running version of GDB, you can run newer commands and syntaxes without breaking legacy scripts.

- **\$_gdb_setting**

\$_gdb_setting_str**\$_gdb_maint_setting****\$_gdb_maint_setting_str**

With these new convenience functions for accessing GDB settings, you can change the logic of user-defined commands depending on current GDB settings.

- **\$_cimag**
\$_creal

With these new convenience functions, you can get the imaginary and real parts of an imaginary number.

- **\$_shell_exitcode**
\$_shell_exitsignal

With these new convenience variables, you can access the exit code or exit status of shell commands executed by the GDB **shell**, **pipe**, and **make** commands.

New and improved commands

- New command options infrastructure has been provided to better support completion, for example, **CMD -[TAB]** will now show a completion of available command options for CMD.
- Command names can now use the . character.

New commands

- **define-prefix**
With this command, you can define your own prefix commands, for example, **abc def** and **abc def ghi** can now be entirely separate commands.
- **| [command] | shell_command** (another name of this command is **pipe**)
This command executes a given *GDB command*, sending its output to the given *shell_command*.
- **with setting [value] [-- command]**
This command temporarily sets the given *setting* to *value* (if specified) and runs the optional *command*, resetting *setting* when complete.

Changed commands

- **help**
apropos

The commands now use new title styling.

- **printf**
eval

The commands can now print C-style and Ada-style string convenience variables without the need for a running process, for example, when debugging core dumps.

- **info sources [-dirname | -basename] [--] [regex]**
New filtering options have been added. Using them, users can limit results to file, directory, or base names matching a given regular expressions.

- **focus**
winheight

+, **-**, **>**, and **<**

These TUI commands are now case-sensitive.

- **backtrace**
The command now supports new options which can override global display preferences (**set backtrace** and **set print** settings). New options include: **-entry-values**, **frame-arguments**, **-raw-frame-arguments**, **-frame-info**, **-past-main**, **-past-entry**, **-full**, **-no-filters**, and **-hide**.

- **frame apply**
tfaas

faas

The commands now support the new **-past-main** and **-past-entry** command options.

- **info types**
The command now support the new **-q** option to disable printing of some header information similar to **info variables** and **info functions**.

- **info variables**
info functions

whereis

The commands now support the new **-n** option to exclude non-debugging symbols from the output.

Settings

- **set may-call-functions**
show may-call-functions

The default value is **on**. These new commands control whether GDB will attempt to call functions in the program during command execution, for example, the **print expression** command.

- **set print finish**
show print finish

If these commands are set to **on**, GDB will print the value returned by the current function when the **finish** command is used.

- **set print max-depth**
show print max-depth

These new commands limit the display of deeply nested structures to the set number of levels. The default limit of nesting levels is 20.

- **set print raw-values**
show print raw-values

These new commands globally override the use of pretty-printers when printing a value. The default value is **off**.

- **set style title foreground**
set style title background

set style title intensity

set style highlight foreground

set style highlight background

set style highlight intensity

With these new commands, you can customize the display styling of the title and highlight styles.

- **maint set worker-threads**
maint show worker-threads

Experimental: if these commands are set to **unlimited**, GDB will use multi-threaded symbol loading for better performance.

- **set style tui-border foreground**
set style tui-border background

set style tui-border intensity

set style tui-active-border foreground

set style tui-active-border background

set style tui-active-border intensity

These new commands set the display styling of various TUI borders.

- **set print frame-info**
show print frame-info

These new commands control what frame information is printed by commands that print a frame, for example, **backtrace**, **frame**, and **stepi**.

- **set tui compact-source**
show tui compact-source

These new commands enable a new **compact** display style for the TUI source window.

- **set debug remote-packet-max-chars**
show debug remote-packet-max-chars

These new commands control the number of characters to output in a remote packet when using **set debug remote**. The default value is 512 bytes.

- **show style**

This new command now styles the output of all subcommands using their own styling.

- **set print frame-arguments**

A new value **presence** has been added. It will display only the presence of arguments with **...** instead of printing the actual argument names and values.

- **set print-raw-frame-arguments**

show print-raw-frame-arguments

These two commands replace deprecated **set/show print raw frame-arguments** commands.

Language-specific improvements

Fortran

- GDB can now set breakpoints on a nested function or subroutine using the **::** operator.
- **info modules [-q] [regex]**
This new command returns a list of modules that match *regex*. If *regex* is not provided, the command returns a list of all modules.
- **info module functions [-q] [-m module_regex] [-t type_regex] [regex]**
info module variables [-q] [-m module_regex] [-t type_regex] [regex]

These new commands return a list of functions or variables within all modules, grouped by module. Results may be limited by module regular expressions, function or variable type signature regular expressions, or name regular expressions.

Python API

- **gdb.Value.format_string**
This new method returns a string representation of the Value object.
- **gdb.Type.objfile**
This new property returns the **objfile** in which the type was defined.
- **gdb.lookup_static_symbol**
gdb.lookup_static_symbols

These new functions support the lookup of symbols with static linkage. The first function returns the first matching symbol. The second one returns all matching symbols.

- **gdb.Objfile.lookup_global_symbol**
gdb.Objfile.lookup_static_symbol

These new functions support lookup of symbols in an **Objfile**: global symbols and symbols with static linkage respectively.

- **gdb.Block**
This function now supports Python dictionary syntax, for example: **symbol = some_block[variable]**, where *symbol* is of type **gdb.Symbol**.

Machine Interface (MI) improvements

- A new default MI version 3 has been introduced (**-i=mi3**).
 - The output of multi-location breakpoint locations has been corrected.
 - The new **-fix-multi-location-breakpoint-output** command has been added to correct syntax errors in older MI versions.
- All new commands are MI implementations of CLI equivalents:
 - **-complete LINE**

- **-catch-throw**
- **-catch-rethrow**
- **-catch-catch**
- **-symbol-info-functions**
- **-symbol-info-types**
- **-symbol-info-variables**
- **-symbol-info-modules**
- **-symbol-info-module-functions**
- **-symbol-info-module-variables**

A.5. CHANGES IN STRACE

Red Hat Developer Toolset 10.0 is distributed with **strace 5.7**.

The following features have been added or modified since the previous release of Red Hat Developer Toolset:

Changes in behavior

- The **%process** class now contains system calls associated with the process life cycle (creation, execution, termination):
 - **kill**, **tkill**, **tgkill**, **pidfd_send_signal**, and **rt_sigqueueinfo** have been added.
 - **arch_prctl** and **unshare** have been removed.
- Messages about unknown tracees are now subject to the **strace** quietness setting **-q** (**--quiet**).
- A new warning has been added. It occurs when the **-A** (**--output-append-mode**) option is used without **-o** (**--output**) or the **-S** (**--summary-sort-by**) option without **-c/-C** (**--summary-only/--summary**).

Improvements

- Every short option now has a long option alias. This change brings the following improvements:
 - The ability to use human-readable settings for the **-I** (**--interruptible**) option.
 - The ability to silence specific messages using the **-e quiet** (**--quiet**) qualifier (which is an alias for the **-q** option), including messages that could not be silenced previously, for example path resolution messages and messages about process being superseded by **execve**.
 - The ability to specify selected flexible data (FD) decoding features using the **-e decode-fds** (**--decode-fds**) qualifier, which is an alias for the **-y** option.
 - The ability to set precision for the absolute timestamp, relative timestamp, and system call time output. Use the **--absolute-timestamps**, **--relative-timestamps**, and **--syscall-times** options, respectively.

- System call return status filtering has been implemented with the **-e status=set** option and its aliases:
 - The **-z (--successful-only)** option limits system call printing to successful system calls only.
 - The **-Z (--failed-only)** option limits system call printing to failed system calls only.
- The **--pidns-translation** option for PID (process ID) namespace translation has been added. This improvement addresses the Fedora bug BZ#1035433.
- Seccomp-BPF can now be used to stop traces only for filtered system calls. Use the **--seccomp-bpf** option to enable this feature.
- Two new extensions to the **-D (--daemonize)** option have been implemented. They move **strace** into a separate process group (**-DD** or **--daemonize=pgroup**) and session (**-DDD** or **--daemonize=session**).
- Interval specification in the **when=** subexpression of the system call tampering expressions has been implemented.
- It is now possible to select the set of displayed columns in the call summary output using the **-U (--summary-columns)** option.
- It is now possible to sort on any summary column.
- System call count statistics have been enhanced: overhead is now applied per call.
- It is now possible to show information about minimum and maximum call duration in the call summary output.
- The system call delay injection and overhead values can now be supplied with a time measure unit suffix and provided in the IEEE 754 floating-point format. Please refer to the “Time specification format description” section of the strace manual page (available via the “scl enable devtoolset-10 – man strace” command) for details.
- Printing of PIDs associated with process file descriptors (pidfds) in **-yy (--decode-fds=pidfd)** mode has been implemented.
- Performance of the **libdw-based** stack traces printing has been improved by implementing a symbol-to-address cache.
- The **-e trace=%clock** option has been added for tracing system calls reading of modifying system clocks.
- The **-e trace=%creds** option has been added for tracing system calls related to process credentials.
- Decoding of the following system calls has been implemented: **clone3**, **fsconfig**, **fsmount**, **fsopen**, **fspick**, **open_tree**, **openat2**, **move_mount**, **pidfd_getfd**, and **pidfd_open**.
- Decoding of the following system calls has been enhanced: **arch_prctl**, **bpf**, **clone**, **inotify_init**, **io_cancel**, **io_submit**, **io_uring_register**, **io_uring_setup**, **keyctl**, **mbind**, **perf_event_open**, **prctl**, **s390_sthyi**, **sched_getattr**, **sched_setattr**, **set_mempolicy**, **syscall**, and **syslog**.
- Decoding of the following **ioctl** commands has been implemented: **PTP_CLOCK_GETCAPS2**, **PTP_EXTTTS_REQUEST2**, **PTP_PEROUT_REQUEST2**, **PTP_ENABLE_PPS2**, **PTP_SYS_OFFSET2**, **RTC_VL_READ**, and **WDIOC_***.

- The **HIDIOCGRAWUNIQL** **ioctl** command number printing has been implemented.
- Decoding of the **NETLINK_ROUTE** netlink protocol has been enhanced.
- Decoding of the following netlink attributes has been implemented: **IFLA_***, **TCA_ACT_FLAGS**, **TCA_STATS_PKT64**, and **UNIX_DIAG_UID**.
- Lists of the following constants have been updated: **AT_***, **AUDIT_***, **BPF_***, **BTRFS_***, **CAN_***, **CLONE_***, **ETH_***, **FAN_***, **GRND_***, **IFLA_***, **IORING_***, **IPPROTO_***, **KEXEC_***, **KEY_***, **KEYCTL_***, **KVM_***, **LWTUNNEL_***, **MADV_***, ***MAGIC**, **MAP_***, **MPOL_***, **MREMAP_***, **MSG_***, **P_***, **PERF_***, **PPC_PTRACE_***, **PR_***, **PTP_***, **RTM_F_***, **SCHED_***, **SCTP_***, **SECCOMP_***, **SO_***, **STATX_***, **TCP_***, **TIPC_***, **UFFDIO_***, **V4L2_***, and **XDP_***.
- The **strace** manual page and the **strace --help** command output have been enhanced.

Bug fixes

- The **statx** system call has been added to the **%fstat** trace class.
- Decoding of **getdents** and **getdents64** system calls has been fixed in cases when they return many directory entries.
- The pathtrace matching of the **openat2** system call has been fixed.
- Various minor fixes in **VIDIOC_*** **ioctl** output formatting have been made.
- Stack trace printing has been fixed for early system calls when **strace** is configured to use **libdw** back end for stack tracing. This improvement addresses the Fedora bug [BZ#1788636](#)
- Decoding of the **NDA_LLADDR** netlink neighbor table attribute has been fixed.
- Decoding of the **BPF_PROG_LOAD** BPF system call command has been fixed.
- The **evdev ioctl** bitset decoding has been fixed.

A.6. CHANGES IN SYSTEMTAP

Red Hat Developer Toolset 10.0 is distributed with **SystemTap 4.3**.

The following features have been added or modified since the previous release of Red Hat Developer Toolset:

- Userspace probes can be targeted by the hexadecimal **buildid** from **readelf -n**. This alternative to a path name enables matching binaries to be probed under any name, and thus allows a single script to target a range of different versions. This feature works well in conjunction with the elfutils **debuginfod** server.
- Script functions can use probe **\$context** variables to access variables in the probed location, which allows the SystemTap scripts to use common logic to work with a variety of probes.

For further information about notable changes, see the upstream [SystemTap 4.3 release notes](#) before updating SystemTap.

A.7. CHANGES IN VALGRIND

Red Hat Developer Toolset 10.0 is distributed with **Valgrind 3.16.1**.

The following features have been added or modified since the previous release of Red Hat Developer Toolset:

- It is now possible to dynamically change the value of many command-line options while your program is running under Valgrind: through **vgdb**, through **gdb** connected to the Valgrind **gdbserver**, or through program client requests. To get a list of dynamically changeable options, run the **valgrind --help-dyn-options** command.
- For the Cachegrind (**cg_annotate**) and Callgrind (**callgrind_annotate**) tools the **--auto** and **--show-percs** options now default to **yes**.
- The Memcheck tool produces fewer false positive errors on optimized code. In particular, Memcheck now better handles the case when the compiler transformed an **A && B** check into **B && A**, where **B** could be undefined and **A** was false. Memcheck also better handles integer equality checks and non-equality checks on partially defined values.
- The experimental Stack and Global Array Checking tool (**exp-sgcheck**) has been removed. An alternative for detecting stack and global array overruns is using the AddressSanitizer (ASAN) facility of GCC, which requires you to rebuild your code with the **-fsanitize=address** option.

APPENDIX B. CHANGES IN VERSION 10.1

The sections below document features and bug fixes introduced in Red Hat Developer Toolset 10.1.

B.1. CHANGES IN GCC

Red Hat Developer Toolset 10.1 is distributed with **GCC 10.2.1**, which provides numerous bug fixes.

One of the most notable bug fixes is:

- Previously, the **gcc-gfortran** subpackage did not include the **libgfortran_nonshared.a** static library. As a consequence, linking a program was not always successful because several symbols were missing. With this update, the **devtoolset-10-gcc-gfortran** package now contains the **libgfortran_nonshared.a** library and the program can be linked successfully.

[Bug#1927579](#)

B.2. CHANGES IN ELFUTILS

Red Hat Developer Toolset 10.1 is distributed with **elfutils 0.182**.

Numerous bug fixes and new features have been added since the previous release of Red Hat Developer Toolset:

- New features added to the **debuginfod** server:
 - More efficient package traversal, **debuginfod** now tolerates various errors during scanning. The grooming process is more visible and interruptible, and provides more Prometheus metrics.
 - A new thread-busy metric and more detailed error metrics.
 - The new **--fdcache-mintmp** option and file system free space tracking.
 - Increased web API concurrency while grooming.
- New features added to **debuginfod-client**:
 - Compressed kernel ELF images are now supported.
 - The **DEBUGINFOD_SONAME** macro has been added to **debuginfod.h**. This macro can be used with the **dlopen()** function to load the **libdebuginfod.so** library.
 - The new **debuginfod_set_verbose_fd** function and **DEBUGINFOD_VERBOSE** environment variable have been added.

B.3. CHANGES IN STRACE

Red Hat Developer Toolset 10.1 is distributed with **strace 5.7**.

The following feature has been added since the previous release of Red Hat Developer Toolset:

- The **--pidns-translation** option has been added for PID namespace translation. For more information about PID namespace translation, see the [pid_namespaces manual page](#). This enhancement addresses [RHEL Bug#1790836](#) and [Fedora Bug#1035433](#).

B.4. CHANGES IN DYNINST

Red Hat Developer Toolset 10.1 is distributed with **Dyninst 10.2.1**.

The following features have been added since the previous release of Red Hat Developer Toolset 10.1:

- support for the **debuginfod** protocol for automated DWARF **debuginfo** downloading
- several memory leak fixes

B.5. CHANGES IN SYSTEMTAP

Red Hat Developer Toolset 10.1 is distributed with **SystemTap 4.4**.

The following features have been added or modified since the previous release of Red Hat Developer Toolset:

- Performance and stability improvements to user-space probing.
- Users can now access implicit thread local storage variables on these architectures: AMD64, Intel 64, IBM Z, the little-endian variant of IBM Power Systems.
- Initial support for processing of floating point values.
- Improved concurrency for scripts using global variables. The locks required to protect concurrent access to global variables have been optimized so that they span the smallest possible critical region.
- New syntax for defining aliases with both a prologue and an epilogue.
- New **@probewrite** predicate.
- **syscall** arguments are writable again.

For further information about notable changes, see the upstream [SystemTap 4.4 release notes](#) before updating SystemTap.