



# Performance Analyzer

---

Sun™ Studio 8

Sun Microsystems, Inc.  
www.sun.com

Part Number 817-5068-10  
March 2004, Revision A

Submit comments about this document at: <http://www.sun.com/hwdocs/feedback>

Copyright © 2004 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, California 95054, U.S.A. All rights reserved.

U.S. Government Rights - Commercial software. Government users are subject to the Sun Microsystems, Inc. standard license agreement and applicable provisions of the FAR and its supplements. Use is subject to license terms.

This distribution may include materials developed by third parties.

Parts of the product may be derived from Berkeley BSD systems, licensed from the University of California. UNIX is a registered trademark in the U.S. and in other countries, exclusively licensed through X/Open Company, Ltd.

Sun, Sun Microsystems, the Sun logo, Java, and JavaHelp are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and other countries. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and other countries. Products bearing SPARC trademarks are based upon architecture developed by Sun Microsystems, Inc.

This product is covered and controlled by U.S. Export Control laws and may be subject to the export or import laws in other countries. Nuclear, missile, chemical biological weapons or nuclear maritime end uses or end users, whether direct or indirect, are strictly prohibited. Export or reexport to countries subject to U.S. embargo or to entities identified on U.S. export exclusion lists, including, but not limited to, the denied persons and specially designated nationals lists is strictly prohibited.

DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.

---

Copyright © 2004 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, California 95054, Etats-Unis. Tous droits réservés.

L'utilisation est soumise aux termes de la Licence.

Cette distribution peut comprendre des composants développés par des tierces parties.

Des parties de ce produit pourront être dérivées des systèmes Berkeley BSD licenciés par l'Université de Californie. UNIX est une marque déposée aux Etats-Unis et dans d'autres pays et licenciée exclusivement par X/Open Company, Ltd.

Sun, Sun Microsystems, le logo Sun, Java, et JavaHelp sont des marques de fabrique ou des marques déposées de Sun Microsystems, Inc. aux Etats-Unis et dans d'autres pays. Toutes les marques SPARC sont utilisées sous licence et sont des marques de fabrique ou des marques déposées de SPARC International, Inc. aux Etats-Unis et dans d'autres pays. Les produits portant les marques SPARC sont basés sur une architecture développée par Sun Microsystems, Inc.

Ce produit est soumis à la législation américaine en matière de contrôle des exportations et peut être soumis à la réglementation en vigueur dans d'autres pays dans le domaine des exportations et importations. Les utilisations, ou utilisateurs finaux, pour des armes nucléaires, des missiles, des armes biologiques et chimiques ou du nucléaire maritime, directement ou indirectement, sont strictement interdites. Les exportations ou réexportations vers les pays sous embargo américain, ou vers des entités figurant sur les listes d'exclusion d'exportation américaines, y compris, mais de manière non exhaustive, la liste de personnes qui font objet d'un ordre de ne pas participer, d'une façon directe ou indirecte, aux exportations des produits ou des services qui sont régis par la législation américaine en matière de contrôle des exportations et la liste de ressortissants spécifiquement désignés, sont rigoureusement interdites.

LA DOCUMENTATION EST FOURNIE "EN L'ÉTAT" ET TOUTES AUTRES CONDITIONS, DECLARATIONS ET GARANTIES EXPRESSES OU TACITES SONT FORMELLEMENT EXCLUES, DANS LA MESURE AUTORISÉE PAR LA LOI APPLICABLE, Y COMPRIS NOTAMMENT TOUTE GARANTIE IMPLICITE RELATIVE A LA QUALITE MARCHANDE, A L'APTITUDE A UNE UTILISATION PARTICULIERE OU A L'ABSENCE DE CONTREFAÇON.



Please  
Recycle



Adobe PostScript™

# Contents

---

<b>Before You Begin</b>	<b>13</b>
How This Book Is Organized	13
Typographic Conventions	14
Shell Prompts	15
Accessing Sun Studio Software and Man Pages	15
Accessing the Compilers and Tools	16
Accessing the Man Pages	17
Accessing the Integrated Development Environment	17
Accessing Compilers and Tools Documentation	18
Documentation in Accessible Formats	19
Related Compilers and Tools Documentation	19
Accessing Related Solaris Documentation	20
Resources for Developers	21
Contacting Sun Technical Support	21
Sending Your Comments	21
<b>1. Overview of the Performance Analyzer</b>	<b>23</b>
Starting the Performance Analyzer From the Integrated Development Environment	23
The Tools of Performance Analysis	23

The Collector Tool	24
The Performance Analyzer Tool	24
The <code>er_print</code> Command	25
The <code>prof</code> , <code>gprof</code> , and <code>tcov</code> Tools	25
The Performance Analyzer Window	25
<b>2. Performance Data</b>	<b>27</b>
Use the Online Help	27
What Data the Collector Collects	28
Clock Data	29
Hardware-Counter Overflow Profiling Data	30
Synchronization Wait Tracing Data	33
Heap Tracing (Memory Allocation) Data	34
MPI Tracing Data	35
Global (Sampling) Data	37
How Metrics Are Assigned to Program Structure	38
Function-Level Metrics: Exclusive, Inclusive, and Attributed	38
Interpreting Attributed Metrics: An Example	39
How Recursion Affects Function-Level Metrics	41
<b>3. Collecting Performance Data</b>	<b>43</b>
Compiling and Linking Your Program	43
Source Code Information	44
Static Linking	44
Optimization	44
Compiling Java Programs	45
Preparing Your Program for Data Collection and Analysis	45
Using Dynamically Allocated Memory	45
Using System Libraries	46

Using Signal Handlers	47
Using <code>setuid</code>	48
Program Control of Data Collection	48
The C, C++, Fortran, and Java API Functions	51
Dynamic Functions and Modules	52
Limitations on Data Collection	53
Limitations on Clock-Based Profiling	54
Limitations on Collection of Tracing Data	54
Limitations on Hardware-Counter Overflow Profiling	55
Runtime Distortion and Dilation With Hardware-Counter Overflow Profiling	56
Limitations on Data Collection for Descendant Processes	56
Limitations on Java Profiling	56
Runtime Performance Distortion and Dilation for Applications Written in the Java Programming Language	57
Where the Data Is Stored	58
Experiment Names	58
Moving Experiments	59
Estimating Storage Requirements	60
Collecting Data	61
Collecting Data Using the <code>collect</code> Command	62
Data Collection Options	62
Experiment Control Options	65
Output Options	68
Other Options	69
Collecting Data Using the <code>dbx collector</code> Subcommands	70
Data Collection Subcommands	71
Experiment Control Subcommands	74
Output Subcommands	74

Information Subcommands	76
Collecting Data From a Running Process	76
Collecting Data From MPI Programs	78
Storing MPI Experiments	79
Running the <code>collect</code> Command Under MPI	81
Collecting Data by Starting <code>dbx</code> Under MPI	81
Using <code>collect</code> With <code>ppgsz</code>	82
<b>4. The <code>er_print</code> Command Line Performance Analysis Tool</b>	<b>83</b>
<code>er_print</code> Syntax	84
Metric Lists	84
Commands That Control the Function List	87
Commands That Control the Callers-Callees List	89
Commands That Control the Leak and Allocation Lists	91
Commands That Control the Source and Disassembly Listings	92
Commands That Control the Data Space List	95
Commands That List Experiments, Samples, Threads, and LWPs	96
Commands That Control Selections	97
Commands That Control Load Object Selection	99
Commands That List Metrics	99
Commands That Control Output	100
Commands That Print Other Displays	101
Commands That Set Defaults	102
Commands That Set Defaults Only For the Performance Analyzer	104
Miscellaneous Commands	105
Examples	106
<b>5. Understanding the Performance Analyzer and Its Data</b>	<b>107</b>
How Data Collection Works	107

Experiment Format	108
Recording Experiments	110
Interpreting Performance Metrics	111
Clock-Based Profiling	111
Synchronization Wait Tracing	114
Hardware-Counter Overflow Profiling	115
Heap Tracing	116
MPI Tracing	116
Call Stacks and Program Execution	117
Single-Threaded Execution and Function Calls	118
Explicit Multithreading	121
Overview of Java Technology-Based Software Execution	122
Java Processing Representations	123
Parallel Execution and Compiler-Generated Body Functions	124
Incomplete Stack Unwinds	128
Mapping Addresses to Program Structure	130
The Process Image	130
Load Objects and Functions	130
Aliased Functions	131
Non-Unique Function Names	131
Static Functions From Stripped Shared Libraries	132
Fortran Alternate Entry Points	132
Cloned Functions	133
Inlined Functions	133
Compiler-Generated Body Functions	134
Outline Functions	135
Dynamically Compiled Functions	135
The <Unknown> Function	135

The <no Java callstack recorded> Function	136
The <Total> Function	136
Mapping Data Addresses to Program Data Objects	137
Dataobject Descriptors	138
Annotated Code Listings	140
Annotated Source Code	140
Annotated Disassembly Code	142
<b>6. Manipulating Experiments and Viewing Annotated Code Listings</b>	<b>149</b>
Manipulating Experiments	149
Viewing Annotated Code Listings With <code>er_src</code>	150
Other Utilities	152
The <code>er_archive</code> Utility	152
The <code>er_export</code> Utility	153
<b>A. Profiling Programs With <code>prof</code>, <code>gprof</code>, and <code>tcov</code></b>	<b>155</b>
Using <code>prof</code> to Generate a Program Profile	156
Using <code>gprof</code> to Generate a Call Graph Profile	158
Using <code>tcov</code> for Statement-Level Analysis	161
Creating <code>tcov</code> Profiled Shared Libraries	164
Locking Files	165
Errors Reported by <code>tcov</code> Runtime Functions	165
Using <code>tcov</code> Enhanced for Statement-Level Analysis	167
Creating Profiled Shared Libraries for <code>tcov</code> Enhanced	168
Locking Files	168
<code>tcov</code> Directories and Environment Variables	169
<b>Index</b>	<b>171</b>



# Figures

---

- [FIGURE 2-1](#) Call Tree Illustrating Exclusive, Inclusive, and Attributed Metrics 40
- [FIGURE 5-1](#) Schematic Call Tree for a Multithreaded Program That Contains a Parallel Do or Parallel For Construct 126
- [FIGURE 5-2](#) Schematic Call Tree for a Parallel Region With a Worksharing Do or Worksharing For Construct 127



# Tables

---

<a href="#">TABLE P-1</a>	Typeface Conventions	14
<a href="#">TABLE P-2</a>	Code Conventions	14
<a href="#">TABLE 2-1</a>	Timing Metrics	29
<a href="#">TABLE 2-2</a>	Aliased Hardware Counters Available on SPARC and IA Hardware	32
<a href="#">TABLE 2-3</a>	Synchronization Wait Tracing Metrics	34
<a href="#">TABLE 2-4</a>	Memory Allocation (Heap Tracing) Metrics	35
<a href="#">TABLE 2-5</a>	MPI Tracing Metrics	36
<a href="#">TABLE 2-6</a>	Classification of MPI Functions Into Send, Receive, Send and Receive, and Other	36
<a href="#">TABLE 3-1</a>	Parameter List for <code>collector_func_load()</code>	53
<a href="#">TABLE 3-2</a>	Environment Variable Settings for Preloading the Library <code>libcollector.so</code>	78
<a href="#">TABLE 4-1</a>	Options for the <code>er_print</code> Command	84
<a href="#">TABLE 4-2</a>	Metric Type Characters	85
<a href="#">TABLE 4-3</a>	Metric Visibility Characters	85
<a href="#">TABLE 4-4</a>	Metric Name Strings	86
<a href="#">TABLE 4-5</a>	Compiler Commentary Message Classes	93
<a href="#">TABLE 4-6</a>	Additional Options for the <code>dcc</code> Command	94
<a href="#">TABLE 4-7</a>	Timeline Display Mode Options	104
<a href="#">TABLE 4-8</a>	Timeline Display Data Types	104
<a href="#">TABLE 5-1</a>	Data Types and Corresponding File Names	108
<a href="#">TABLE 5-2</a>	How Kernel Microstates Contribute to Metrics	112

TABLE 5-3	Annotated Source-Code Metrics	141
TABLE A-1	Performance Profiling Tools	155

# Before You Begin

---

This manual describes the performance analysis tools that are available with Sun™ Studio 8.

- The Collector and Performance Analyzer are a pair of tools that perform statistical profiling of a wide range of performance data and tracing of various system calls, and relate the data to program structure at the function, source line and instruction level.
- `prof` and `gprof` are tools that perform statistical profiling of CPU usage and provide execution frequencies at the function level.
- `tcov` is a tool that provides execution frequencies at the function and source line levels.

This manual is intended for application developers with a working knowledge of Fortran, C, C++, or Java™, the Solaris™ Operating System, and UNIX® operating system commands. Some knowledge of performance analysis is helpful but is not required to use the tools.

---

## How This Book Is Organized

[Chapter 1](#) introduces the performance analysis tools, briefly discussing what they do and when to use them.

[Chapter 2](#) describes the data collected by the Collector and how the data is converted into metrics of performance.

[Chapter 3](#) describes how to use the Collector to collect timing data, synchronization delay data, and hardware event data from your program.

[Chapter 4](#) describes how to use the `er_print` command line interface to analyze the data collected by the Collector.

[Chapter 5](#) describes the process of converting the data collected by the Collector into performance metrics and how the metrics are related to program structure.

[Chapter 6](#) presents information on the utilities that are provided for manipulating and converting performance experiments and viewing annotated source code and disassembly code without running an experiment.

[Appendix A](#) describes the UNIX profiling tools `prof`, `gprof`, and `tcov`. These tools provide timing information and execution frequency statistics.

---

## Typographic Conventions

**TABLE P-1** Typeface Conventions

Typeface	Meaning	Examples
AaBbCc123	The names of commands, files, and directories; on-screen computer output	Edit your <code>.login</code> file. Use <code>ls -a</code> to list all files. % You have mail.
<b>AaBbCc123</b>	What you type, when contrasted with on-screen computer output	% <b>su</b> Password:
<i>AaBbCc123</i>	Book titles, new words or terms, words to be emphasized	Read Chapter 6 in the <i>User's Guide</i> . These are called <i>class</i> options. You <i>must</i> be superuser to do this.
<code>AaBbCc123</code>	Command-line placeholder text; replace with a real name or value	To delete a file, type <code>rm filename</code> .

**TABLE P-2** Code Conventions

Code Symbol	Meaning	Notation	Code Example
[ ]	Brackets contain arguments that are optional.	<code>O[n]</code>	<code>O4, O</code>
{ }	Braces contain a set of choices for a required option.	<code>d{y n}</code>	<code>dy</code>

**TABLE P-2** Code Conventions (*Continued*)

Code Symbol	Meaning	Notation	Code Example
	The “pipe” or “bar” symbol separates arguments, only one of which may be chosen.	B{dynamic static}	Bstatic
:	The colon, like the comma, is sometimes used to separate arguments.	Rdir[:dir]	R/local/libs:/U/a
...	The ellipsis indicates omission in a series.	xinline= <i>fn</i> [,... <i>fn</i> ]	xinline=alpha,dos

---

## Shell Prompts

Shell	Prompt
C shell	<i>machine-name</i> %
C shell superuser	<i>machine-name</i> #
Bourne shell and Korn shell	\$
Superuser for Bourne shell and Korn shell	#

---

## Accessing Sun Studio Software and Man Pages

The compilers and tools and their man pages are not installed into the standard `/usr/bin/` and `/usr/share/man` directories. To access the compilers and tools, you must have your `PATH` environment variable set correctly (see [“Accessing the Compilers and Tools” on page 16](#)). To access the man pages, you must have the your `MANPATH` environment variable set correctly (see [“Accessing the Man Pages” on page 17](#)).

For more information about the `PATH` variable, see the `csh(1)`, `sh(1)`, and `ksh(1)` man pages. For more information about the `MANPATH` variable, see the `man(1)` man page. For more information about setting your `PATH` variable and `MANPATH` variables to access this release, see the installation guide or your system administrator.

---

**Note** – The information in this section assumes that your Sun Studio compilers and tools are installed in the `/opt` directory. If your software is not installed in the `/opt` directory, ask your system administrator for the equivalent path on your system.

---

## Accessing the Compilers and Tools

Use the steps below to determine whether you need to change your `PATH` variable to access the compilers and tools.

### ▼ To Determine Whether You Need to Set Your `PATH` Environment Variable

1. **Display the current value of the `PATH` variable by typing the following at a command prompt.**

```
% echo $PATH
```

2. **Review the output to find a string of paths that contain `/opt/SUNWspro/bin/`.**

If you find the path, your `PATH` variable is already set to access the compilers and tools. If you do not find the path, set your `PATH` environment variable by following the instructions in the next procedure.

### ▼ To Set Your `PATH` Environment Variable to Enable Access to the Compilers and Tools

1. **If you are using the C shell, edit your home `.cshrc` file. If you are using the Bourne shell or Korn shell, edit your home `.profile` file.**
2. **Add the following to your `PATH` environment variable. If you have Sun ONE Studio software or Forte Developer software installed, add the following path before the paths to those installations.**

```
/opt/SUNWspro/bin
```



## Accessing the Man Pages

Use the following steps to determine whether you need to change your `MANPATH` variable to access the man pages.

### ▼ To Determine Whether You Need to Set Your `MANPATH` Environment Variable

1. Request the `dbx` man page by typing the following at a command prompt.

```
% man dbx
```

2. Review the output, if any.

If the `dbx(1)` man page cannot be found or if the man page displayed is not for the current version of the software installed, follow the instructions in the next procedure for setting your `MANPATH` environment variable.

### ▼ To Set Your `MANPATH` Environment Variable to Enable Access to the Man Pages

1. If you are using the C shell, edit your home `.cshrc` file. If you are using the Bourne shell or Korn shell, edit your home `.profile` file.
2. Add the following to your `MANPATH` environment variable.

```
/opt/SUNWspro/man
```

## Accessing the Integrated Development Environment

The Sun Studio 8 integrated development environment (IDE) provides modules for creating, editing, building, debugging, and analyzing the performance of a C, C++, or Fortran application.

The IDE requires the Core Platform component of Sun Studio 8. You must set the `SPRO_NETBEANS_HOME` environment variable to the location where the Core Platform component is installed or mounted (*installation\_directory/netbeans/3.5R*) if the Core Platform component is not installed or mounted to one of the following locations:

- The default installation directory `/opt/netbeans/3.5R`

- The same location as the Compilers and Tools component of the Sun Studio 8 (for example, the Compilers and Tools component installed in `/foo/SUNWspro` and the Core Platform component in `/foo/netbeans/3.5R`)

The command to start the IDE is `sunstudio`. For details on this command, see the `sunstudio(1)` man page.

---

## Accessing Compilers and Tools Documentation

You can access the documentation at the following locations:

- The documentation is available from the documentation index that is installed with the software on your local system or network at `file:/opt/SUNWspro/docs/index.html`.

If your software is not installed in the `/opt` directory, ask your system administrator for the equivalent path on your system.

- Most manuals are available from the `docs.sun.com`<sup>sm</sup> web site. The following titles are available through your installed software only:
  - *Standard C++ Library Class Reference*
  - *Standard C++ Library User's Guide*
  - *Tools.h++ Class Library Reference*
  - *Tools.h++ User's Guide*
- The release notes are available from the `docs.sun.com` web site.
- Online help for all components of the IDE is available through the Help menu, as well as through Help buttons on many windows and dialogs, in the IDE.

The `docs.sun.com` web site (<http://docs.sun.com>) enables you to read, print, and buy Sun Microsystems manuals through the Internet. If you cannot find a manual, see the documentation index that is installed with the software on your local system or network.

---

**Note** – Sun is not responsible for the availability of third-party web sites mentioned in this document and does not endorse and is not responsible or liable for any content, advertising, products, or other materials on or available from such sites or resources. Sun will not be responsible or liable for any damage or loss caused or alleged to be caused by or in connection with use of or reliance on any such content, goods, or services available on or through any such sites or resources.

---

# Documentation in Accessible Formats

The documentation is provided in accessible formats that are readable by assistive technologies for users with disabilities. You can find accessible versions of documentation as described in the following table. If your software is not installed in the `/opt` directory, ask your system administrator for the equivalent path on your system.

Type of Documentation	Format and Location of Accessible Version
Manuals (except third-party manuals)	HTML at <a href="http://docs.sun.com">http://docs.sun.com</a>
Third-party manuals: <ul style="list-style-type: none"><li>• <i>Standard C++ Library Class Reference</i></li><li>• <i>Standard C++ Library User's Guide</i></li><li>• <i>Tools.h++ Class Library Reference</i></li><li>• <i>Tools.h++ User's Guide</i></li></ul>	HTML in the installed software through the documentation index at <code>file:/opt/SUNWspro/docs/index.html</code>
Readmes and man pages	HTML in the installed software through the documentation index at <code>file:/opt/SUNWspro/docs/index.html</code>
Online help	HTML available through the Help menu in the IDE
Release notes	HTML at <a href="http://docs.sun.com">http://docs.sun.com</a>

## Related Compilers and Tools Documentation

The following table describes related documentation that is available at `file:/opt/SUNWspro/docs/index.html` and <http://docs.sun.com>. If your software is not installed in the `/opt` directory, ask your system administrator for the equivalent path on your system.

Document Title	Description
<i>OpenMP API User's Guide</i>	Information on compiler directives used to parallelize programs.

Document Title	Description
<i>Fortran Programming Guide</i>	Discusses programming techniques, including parallelization, optimization, creation of shared libraries.
<i>Debugging a Program With dbx</i>	Reference manual for use of the debugger. Provides information on attaching and detaching to Solaris processes, and executing programs in a controlled environment.
Language user's guides	Describe compilation and compiler options.

## Accessing Related Solaris Documentation

The following table describes related documentation that is available through the docs.sun.com web site.

Document Collection	Document Title	Description
Solaris Reference Manual Collection	See the titles of man page sections.	Provides information about the Solaris Operating System.
Solaris Software Developer Collection	<i>Linker and Libraries Guide</i>	Describes the operations of the Solaris link-editor and runtime linker.
Solaris Software Developer Collection	<i>Multithreaded Programming Guide</i>	Covers the POSIX and Solaris threads APIs, programming with synchronization objects, compiling multithreaded programs, and finding tools for multithreaded programs.
Solaris Software Developer Collection	<i>SPARC Assembly Language Reference Manual</i>	Describes the assembly language for SPARC® processors.
Solaris 9 Update Collection	<i>Solaris Tunable Parameters Reference Manual</i>	Provides reference information on Solaris tunable parameters.

---

## Resources for Developers

Visit <http://developers.sun.com/prodtech/cc> to find these frequently updated resources:

- Articles on programming techniques and best practices
- A knowledge base of short programming tips
- Documentation of compilers and tools components, as well as corrections to the documentation that is installed with your software
- Information on support levels
- User forums
- Downloadable code samples
- New technology previews

You can find additional resources for developers at <http://developers.sun.com>.

---

## Contacting Sun Technical Support

If you have technical questions about this product that are not answered in this document, go to:

<http://www.sun.com/service/contacting>

---

## Sending Your Comments

Sun is interested in improving its documentation and welcomes your comments and suggestions. Email your comments to Sun at this address:

`docfeedback@sun.com`

Please include the part number (817-5068-10) of your document in the subject line of your email.



# Overview of the Performance Analyzer

---

Developing high performance applications requires a combination of compiler features, libraries of optimized functions, and tools for performance analysis. The *Performance Analyzer* manual describes the tools that are available to help you assess the performance of your code, identify potential performance problems, and locate the part of the code where the problems occur.

---

## Starting the Performance Analyzer From the Integrated Development Environment

For information on starting the Performance Analyzer from the Integrated Development Environment (IDE), see the Performance Analyzer Readme, which is available through the documentation index at the default installation directory: `file:/opt/SUNWspro/docs/index.html`. If the Sun Studio 8 compilers and tools are not installed in the `/opt` directory, ask your system administrator for the equivalent path on your system.

---

## The Tools of Performance Analysis

This manual deals primarily with the Collector and Performance Analyzer, a pair of tools that you use to collect and analyze performance data for your application. Both tools can be used from the command line or from a graphical user interface.

The Collector and Performance Analyzer are designed for use by any software developer, even if performance tuning is not the developer's main responsibility. These tools provide a more flexible, detailed, and accurate analysis than the commonly used profiling tools `prof` and `gprof`, and are not subject to an attribution error in `gprof`.

These two tools help to answer the following kinds of questions:

- How much of the available resources does the program consume?
- Which functions or load objects are consuming the most resources?
- Which source lines and instructions are responsible for resource consumption?
- How did the program arrive at this point in the execution?
- Which resources are being consumed by a function or load object?

## The Collector Tool

The Collector tool collects performance data using a statistical method called profiling and by tracing function calls. The data can include call stacks, microstate accounting information, thread-synchronization delay data, hardware-counter overflow data, Message Passing Interface (MPI) function call data, memory allocation data and summary information for the operating system and the process. The Collector can collect all kinds of data for C, C++ and Fortran programs, and it can collect profiling data for applications written in the Java™ programming language. It can collect data for dynamically-generated functions and for descendant processes. See [Chapter 2](#) for information about the data collected and [Chapter 3](#) for detailed information about the Collector. The Collector can be run from the Performance Analyzer GUI, from the IDE, from the `dbx` command line tool, and using the `collect` command.

## The Performance Analyzer Tool

The Performance Analyzer tool displays the data recorded by the Collector, so that you can examine the information. The Performance Analyzer processes the data and displays various metrics of performance at the level of the program, the functions, the source lines, and the instructions. These metrics are classed into five groups:

- Timing metrics
- Hardware counter metrics
- Synchronization delay metrics
- Memory allocation metrics
- MPI tracing metrics.

The Performance Analyzer also displays the raw data in a graphical format as a function of time. The Performance Analyzer can create a mapfile that you can use to improve the order of function loading in the program's address space.



See the online help in the IDE or the Performance Analyzer GUI for detailed information about the Performance Analyzer, and [Chapter 4](#) for information about the command-line analysis tool, `er_print`.

[Chapter 5](#) discusses topics related to understanding the performance analyzer and its data, including: how data collection works, interpreting performance metrics, call stacks and program execution, and annotated code listings. Annotated source code listings and disassembly code listings that include compiler commentary but do not include performance data can be viewed with the `er_src` utility (see [Chapter 6](#) for more information).

## The `er_print` Command

The `er_print` command presents in plain text all the displays that are presented by the Performance Analyzer, with the exception of the Timeline display.

## The `prof`, `gprof`, and `tcov` Tools

This manual also includes information about the following performance tools:

- `prof` and `gprof`

`prof` and `gprof` are UNIX® tools for generating profile data and are included with the Solaris 7, Solaris 8, and Solaris 9 Operating System (SPARC® Platform Edition). Both tools are also provided and supported on the x86 platform.

- `tcov`

`tcov` is a code coverage tool that reports the number of times each function is called and each source line is executed.

For more information about `prof`, `gprof`, and `tcov`, see [Appendix A](#).

---

# The Performance Analyzer Window

---

**Note** – The following is a brief overview of the Performance Analyzer window. See the online help for a complete and detailed discussion of the functionality and features of the tabs discussed below.

---

The Performance Analyzer window consists of a multi-tabbed display, with a menu bar and a toolbar. The tab that is displayed when the Performance Analyzer is started shows a list of functions for the program with exclusive and inclusive metrics for each function. The list can be filtered by load object, by thread, by Light Weight Process (LWP), and by time slice.

For a selected function, another tab displays the callers and callees of the function. This tab can be used to navigate the call tree—in search of high metric values, for example.

Two other tabs display source code that is annotated line-by-line with performance metrics and interleaved with compiler commentary, and disassembly code that is annotated with metrics for each instruction and interleaved with both source code and compiler commentary if they are available.

The performance data is displayed as a function of time in another tab.

Other tabs show details of the experiments and load objects, summary information for a function, and statistics for the process.

You can navigate the Performance Analyzer from the keyboard as well as a mouse.

## Performance Data

---

The performance tools work by recording data about specific events while a program is running, and converting the data into measurements of program performance called metrics.

This chapter describes the data collected by the performance tools, how it is processed and displayed, and how it can be used for performance analysis. Because there is more than one tool that collects performance data, the term Collector is used to refer to any of these tools. Likewise, because there is more than one tool that analyzes performance data, the term analysis tools is use to refer to any of these tools.

This chapter covers the following topics.

- [Use the Online Help](#)
- [What Data the Collector Collects](#)
- [How Metrics Are Assigned to Program Structure](#)

See [Chapter 3](#) for information on collecting and storing performance data.

See [Chapter 4](#) for information on analyzing performance data with `er_print`.

---

## Use the Online Help

For more information analyzing performance data, see the online help in the IDE or the online help in the Performance Analyzer GUI.

---

# What Data the Collector Collects

The Collector collects three different kinds of data: profiling data, tracing data and global data.

- Profiling data is collected by recording profile events at regular intervals. The interval is either a time interval obtained by using the system clock or a number of hardware events of a specific type. When the interval expires, a signal is delivered to the system and the data is recorded at the next opportunity.
- Tracing data is collected by interposing a wrapper function on various system functions so that calls to the system functions can be intercepted and data recorded about the calls.
- Global data is collected by calling various system routines to obtain information. The global data packet is called a sample.

Both profiling data and tracing data contain information about specific events, and both types of data are converted into performance metrics. Global data is not converted into metrics, but is used to provide markers that can be used to divide the program execution into time segments. The global data gives an overview of the program execution during that time segment.

The data packets collected at each profiling event or tracing event include the following information:

- A header identifying the data
- A high-resolution timestamp
- A thread ID
- A lightweight process (LWP) ID
- A processor ID, when available from the operating system
- A copy of the call stack

For more information on threads and lightweight processes, see [Chapter 5](#).

In addition to the common data, each event-specific data packet contains information specific to the data type. The five types of data that the Collector can record are:

- Clock profile data
- Hardware-counter overflow profiling data
- Synchronization wait tracing data
- Heap tracing (memory allocation) data
- MPI tracing data

These five data types, the metrics that are derived from them, and how you might use them, are described in the next five subsections.

# Clock Data

In clock-based profiling, the state of each LWP is stored at regular time intervals. This time interval is called the profiling interval. The information is stored in an integer array: one element of the array is used for each of the ten microaccounting states maintained by the kernel. The data collected is converted by the Performance Analyzer into times spent in each state, with a resolution of the profiling interval. The default profiling interval is approximately 10 ms. The Collector provides a high-resolution profiling interval of approximately 1 ms and a low-resolution profiling interval of approximately 100 ms., and, where the OS permits, allows arbitrary intervals. Running `collect` with no arguments will print the range and resolution allowable on the system on which it is run.

The metrics that are computed from clock-based data are defined in the following table.

**TABLE 2-1** Timing Metrics

Metric	Definition
User CPU time	LWP time spent running in user mode on the CPU.
Wall time	LWP time spent in LWP 1. This is the “wall clock time”
Total LWP time	Sum of all LWP times.
System CPU time	LWP time spent running in kernel mode on the CPU or in a trap state.
Wait CPU time	LWP time spent waiting for a CPU.
User lock time	LWP time spent waiting for a lock.
Text page fault time	LWP time spent waiting for a text page.
Data page fault time	LWP time spent waiting for a data page.
Other wait time	LWP time spent waiting for a kernel page, or time spent sleeping or stopped.

For multithreaded experiments, times other than wall clock time are summed across all LWPs. Wall time as defined is not meaningful for multiple-program multiple-data (MPMD) programs.

Timing metrics tell you where your program spent time in several categories and can be used to improve the performance of your program.

- High user CPU time tells you where the program did most of the work. It can be used to find the parts of the program where there may be the most gain from redesigning the algorithm.
- High system CPU time tells you that your program is spending a lot of time in calls to system routines.

- High wait CPU time tells you that there are more threads ready to run than there are CPUs available, or that other processes are using the CPUs.
- High user lock time tells you that threads are unable to obtain the lock that they request.
- High text page fault time means that the code generated by the linker is organized in memory so that calls or branches cause a new page to be loaded. Creating and using a mapfile (see “Generating and Using a Mapfile” in the Performance Analyzer online help) can fix this kind of problem.
- High data page fault time indicates that access to the data is causing new pages to be loaded. Reorganizing the data structure or the algorithm in your program can fix this problem.

## Hardware-Counter Overflow Profiling Data

Hardware counters keep track of events like cache misses, cache stall cycles, floating-point operations, branch mispredictions, CPU cycles, and instructions executed. In hardware-counter overflow profiling, the Collector records a profile packet when a designated hardware counter of the CPU on which an LWP is running overflows. The counter is reset and continues counting. The profile packet includes the overflow value and the counter type.

The UltraSPARC® III processor family and the IA processor family have two registers that can be used to count events. The Collector can collect data from either or both registers. For each register the Collector allows you to select the type of counter to monitor for overflow, and to set an overflow value for the counter. Some hardware counters can use either register, others are only available on a particular register. Consequently, not all combinations of hardware counters can be chosen in a single experiment.

Hardware-counter overflow profiling data is converted by the Performance Analyzer into count metrics. For counters that count in cycles, the metrics reported are converted to times; for counters that do not count in cycles, the metrics reported are event counts. On machines with multiple CPUs, the clock frequency used to convert the metrics is the harmonic mean of the clock frequencies of the individual CPUs. Because each type of processor has its own set of hardware counters, and because the number of hardware counters is large, the hardware counter metrics are not listed here. The next subsection tells you how to find out what hardware counters are available.

One use of hardware counters is to diagnose problems with the flow of information into and out of the CPU. High counts of cache misses, for example, indicate that restructuring your program to improve data or text locality or to increase cache reuse can improve program performance.

Some of the hardware counters provide similar or related information. For example, branch mispredictions and instruction cache misses are often related because a branch misprediction causes the wrong instructions to be loaded into the instruction cache, and these must be replaced by the correct instructions. The replacement can cause an instruction cache miss, or an instruction translation lookaside buffer (ITLB) miss.

Hardware-counter overflows are often delivered one or more instructions after the instruction which caused the event and the corresponding event counter to overflow: this is referred to as “skid” and it can make counter overflow profiles difficult to interpret. In the absence of hardware support for precise identification of the causal instruction, an apropos backtracking search for a candidate causal instruction may be attempted.

When such backtracking is supported and specified during collection, hardware counter profile packets additionally include the PC (program counter) and EA (effective address) of a candidate memory-referencing instruction appropriate for the hardware counter event. (Subsequent processing during analysis is required to validate the candidate event PC and EA.) This additional information about memory-referencing events facilitates various data-oriented analyses.

## Hardware Counter Lists

Hardware counters are processor-specific, so the choice of counters available to you depends on the processor that you are using. The performance tools provide aliases for a number of counters that are likely to be in common use. You can obtain a list of available hardware counters on any particular system from the Collector by typing `collect` with no arguments in a terminal window on that system.

Here is an example that shows the entries in the counter list for aliased counters. Each line of output in this example is formatted for print. The actual output does not break the lines.

```
CPU Cycles (cycles = Cycle_cnt/*) 9999991 hi=1000003, lo=10000007
(CPU-cycles)

Instructions Executed (insts = Instr_cnt/*) 9999991 hi=1000003,
lo=100000007 (Events)

D$ Read Misses (dcrm = DC_rd_miss/1) 100003 hi=10007, lo=100003
load (Events)
```

In the first line, the first field, `CPU Cycles`, is the metric name. The second field, `cycles`, gives the alias name that can be used in the `-h counter...` argument of `er_print`. The third field, `Cycle_cnt/*`, gives the internal name as used by `cputrack(1)` and the register number on which that counter can be used. The

register number is either 0 or 1, or \*. In this example, the \* indicates that the counter is available on either register. The next field is the overflow interval, the following field is the high-resolution overflow interval, and the last field is the low-resolution overflow interval. The (CPU-cycles) indicates that the counter counts in units of CPU-cycles, and can be converted to time.

The second line terminates with (Events) which indicates that it counts events, and can not be converted to time.

A line may have an additional field, as shown in the third line of the example, between the low-resolution value and the counter units, indicating whether the counter may be triggered by a load, a store, or either, or if the counter is not program related.

The aliased counters that are available on both UltraSPARC and IA hardware are given in TABLE 2-2. There are other aliases that are available on UltraSPARC hardware.

**TABLE 2-2** Aliased Hardware Counters Available on SPARC and IA Hardware

Aliased Counter Name	Metric Name	Description
cycles	CPU Cycles	CPU cycles, counted on either register
insts	Instructions Executed	Instructions executed, counted on either register

Each line of output in this example is formatted for print. The actual output does not break the lines. Lines of output for the non-aliased counters are as follows:

```
Cycle_cnt Events (reg. 0) 1000003 hi=100003, lo=9999991
(CPU-cycles)

Instr_cnt Events (reg. 0) 1000003 hi=100003, lo=9999991 (Events)

DC_rd Events (reg. 0) 1000003 hi=100003, lo=9999991 load (Events)
```

In this line, the first field, Cycle\_cnt, gives the internal name as used by cputrack(1) and the register number on which that counter can be used. The string Cycle\_cnt Events is the metric name for this counter. The remainder of the line is formatted as for an aliased counter.

For both aliased and non-aliased counters that count in cycles, indicated by a (CPU-cycles) at the end of its line, the metrics reported are converted by default to inclusive and exclusive times, but can optionally be shown as event counts. For counters that count in events, indicated by (Events) at the end of its line, the metrics reported are inclusive and exclusive event counts.



For hardware counters that relate to memory operations, as indicated by `load`, `store`, or `load-store` following the counter name, the name of the counter may be preceded by a `+` sign, to request that the data collection attempt to find the precise instruction and effective address that caused the event on the counter that overflowed.

If a counter is not program related, using it for profiling will generate a warning, and profiling will not record a callstack, but rather will show the time being spent in an artificial function, `collector_not_program_related`. Thread and LWP ID's will be recorded, but are meaningless. The string `not-program-related` will appear after the name of any counter that counts events unrelated to the program running. Using such a counter reports metrics in the function `collector_not_program_related` and will give a warning before collection.

In the counter list, the aliased counters appear first, then all the counters available on register 0, then all the counters available on register 1. The aliased counters appear twice, with and without the alias. In the non-aliased list, these counters can have different overflow values. The default overflow values for the aliased cycle counters have been chosen to produce approximately the same data collection rate as for clock data. Other counters are much more sensitive to the actual behavior of the application.

## Synchronization Wait Tracing Data

In multithreaded programs, the synchronization of tasks performed by different threads can cause delays in execution of your program, because one thread might have to wait for access to data that has been locked by another thread, for example. These events are called synchronization delay events and are collected by tracing calls to the functions in the threads library, `libthread.so`. The process of collecting and recording these events is called synchronization wait tracing. The time spent waiting for the lock is called the wait time.

Events are only recorded if their wait time exceeds a threshold value, which is given in microseconds. A threshold value of 0 means that all synchronization delay events are traced, regardless of wait time. The default threshold is determined by running a calibration test, in which calls are made to the threads library without any synchronization delay. The threshold is the average time for these calls multiplied by an arbitrary factor (currently 6). This procedure prevents the recording of events for which the wait times are due only to the call itself and not to a real delay. As a result, the amount of data is greatly reduced, but the count of synchronization events can be significantly underestimated.

Synchronization tracing for Java programs is based on events generated when a thread attempts to acquire a Java Monitor. Both machine and Java callstacks are collected for these events, but no synchronization tracing data is collected for

internal locks used within the JVM. In the machine representation, thread synchronization devolves into calls to `_lwp_mutex_lock`, and no synchronization data is shown, since these calls are not traced.

Synchronization wait tracing data is converted into the following metrics:

**TABLE 2-3** Synchronization Wait Tracing Metrics

Metric	Definition
Synchronization delay events.	The number of calls to a synchronization routine where the wait time exceeded the prescribed threshold.
Synchronization wait time.	Total of wait times that exceeded the prescribed threshold.

From this information you can determine if functions or load objects are either frequently blocked, or experience unusually long wait times when they do make a call to a synchronization routine. High synchronization wait times indicate contention among threads. You can reduce the contention by redesigning your algorithms, particularly restructuring your locks so that they cover only the data for each thread that needs to be locked.

## Heap Tracing (Memory Allocation) Data

Calls to memory allocation and deallocation functions that are not properly managed can be a source of inefficient data usage and can result in poor program performance. In heap tracing, the Collector traces memory allocation and deallocation requests by interposing on the C standard library memory allocation functions `malloc`, `realloc`, `valloc`, and `memalign` and the deallocation function `free`. The Fortran functions `allocate` and `deallocate` call the C standard library functions, so these routines are also traced indirectly.

For Java programs, heap tracing data records all object allocation events (generated by the user code), and object deallocation events (generated by the garbage collector). In addition, any use of `malloc`, `free`, *etc.* also generates events that are recorded. Those events may come from native code, or from the JVM itself. In the machine representation, memory is allocated and deallocated by the JVM, typically in very large chunks. Memory allocation from the Java code is handled entirely by the JVM and its garbage-collector. Heap tracing will not show JVM allocations, since they are done by mapping memory, rather than calling the normal heap routines, and will not show any information about the Java memory allocation and garbage collection.

Heap tracing data is converted into the following metrics:

**TABLE 2-4** Memory Allocation (Heap Tracing) Metrics

<b>Metric</b>	<b>Definition</b>
Allocations	The number of calls to the memory allocation functions.
Bytes allocated	The sum of the number of bytes allocated in each call to the memory allocation functions.
Leaks	The number of calls to the memory allocation functions that did not have a corresponding call to a deallocation function.
Bytes leaked	The number of bytes that were allocated but not deallocated.

Collecting heap tracing data can help you identify memory leaks in your program or locate places where there is inefficient allocation of memory.

There is another definition of memory leaks that is commonly used, such as in the debugging tool, `dbx`. This alternate definition says a memory leak is a dynamically-allocated block of memory that has no pointers pointing to it anywhere in the data space of the program. The definition of leaks used here includes this alternative definition, but also includes memory for which pointers do exist.

## MPI Tracing Data

The Collector can collect data on calls to the Message Passing Interface (MPI) library. The functions for which data is collected are listed below.

<code>MPI_Allgather</code>	<code>MPI_Allgatherv</code>	<code>MPI_Allreduce</code>
<code>MPI_Alltoall</code>	<code>MPI_Alltoallv</code>	<code>MPI_Barrier</code>
<code>MPI_Bcast</code>	<code>MPI_Bsend</code>	<code>MPI_Gather</code>
<code>MPI_Gatherv</code>	<code>MPI_Irecv</code>	<code>MPI_Isend</code>
<code>MPI_Recv</code>	<code>MPI_Reduce</code>	<code>MPI_Reduce_scatter</code>
<code>MPI_Rsend</code>	<code>MPI_Scan</code>	<code>MPI_Scatter</code>
<code>MPI_Scatterv</code>	<code>MPI_Send</code>	<code>MPI_Sendrecv</code>
<code>MPI_Sendrecv_replace</code>	<code>MPI_Ssend</code>	<code>MPI_Wait</code>
<code>MPI_Waitall</code>	<code>MPI_Waitany</code>	<code>MPI_Waitsome</code>
<code>MPI_Win_fence</code>	<code>MPI_Win_lock</code>	

MPI tracing data is converted into the following metrics:

**TABLE 2-5** MPI Tracing Metrics

<b>Metric</b>	<b>Definition</b>
MPI Receives	Number of receive operations in MPI functions that receive data
MPI Bytes Received	Number of bytes received in MPI functions
MPI Sends	Number of send operations in MPI functions that send data
MPI Bytes Sent	Number of bytes sent in MPI functions
MPI Time	Time spent in all calls to MPI functions
Other MPI Calls	Number of calls to other MPI functions

The number of bytes recorded as received or sent is the buffer size given in the call. This might be larger than the actual number of bytes received or sent. In the global communication functions and collective communication functions, the number of bytes sent or received is the maximum number, assuming direct interprocessor communication and no optimization of the data transfer or re-transmission of the data.

The functions from the MPI library that are traced are listed in [TABLE 2-6](#), categorized as MPI send functions, MPI receive functions, MPI send and receive functions, and other MPI functions.

**TABLE 2-6** Classification of MPI Functions Into Send, Receive, Send and Receive, and Other

<b>Category</b>	<b>Functions</b>
MPI send functions	MPI_Bsend, MPI_Isend, MPI_Rsend, MPI_Send, MPI_Ssend
MPI receive functions	MPI_Irecv, MPI_Recv
MPI send and receive functions	MPI_Allgather, MPI_Allgatherv, MPI_Allreduce, MPI_Alltoall, MPI_Alltoallv, MPI_Bcast, MPI_Gather, MPI_Gatherv, MPI_Reduce, MPI_Reduce_scatter, MPI_Scan, MPI_Scatter, MPI_Scatterv, MPI_Sendrecv, MPI_Sendrecv_replace
Other MPI functions	MPI_Barrier, MPI_Wait, MPI_Waitall, MPI_Waitany, MPI_Waitsome, MPI_Win_fence, MPI_Win_lock

Collecting MPI tracing data can help you identify places where you have a performance problem in an MPI program that could be due to MPI calls. Examples of possible performance problems are load balancing, synchronization delays, and communications bottlenecks.

## Global (Sampling) Data

Global data is recorded by the Collector in packets called sample packets. Each packet contains a header, a timestamp, execution statistics from the kernel such as page fault and I/O data, context switches, and a variety of page residency (working-set and paging) statistics. The data recorded in sample packets is global to the program and is not converted into performance metrics. The process of recording sample packets is called sampling.

Sample packets are recorded in the following circumstances:

- When the program stops for any reason in the Debugging window or in `dbx`, such as at a breakpoint, if the option to do this is set
- At the end of a sampling interval, if you have selected periodic sampling. The sampling interval is specified as an integer in units of seconds. The default value is 1 second
- When you choose Debug → Performance Toolkit → New Sample, or click the New Sample button in the Debugging window, or use the `dbx collector sample record` command
- At a call to `collector_sample`, if you have put calls to this routine in your code (see “Program Control of Data Collection” on page 48)
- When a specified signal is delivered, if you have used the `-l` option with the `collect` command (see the `collect(1)` man page)
- When collection is initiated and terminated
- Before and after a descendant process is created

The performance tools use the data recorded in the sample packets to group the data into time periods, which are called samples. You can filter the event-specific data by selecting a set of samples, so that you see only information on a particular time period. You can also view the global data for each sample.

The performance tools make no distinction between the different kinds of sample points. To make use of sample points for analysis you should choose only one kind of point to be recorded. In particular, if you want to record sample points that are related to the program structure or execution sequence, you should turn off periodic sampling, and use samples recorded when `dbx` stops the process, or when a signal is delivered to the process that is recording data using the `collect` command, or when a call is made to the Collector API functions.

---

# How Metrics Are Assigned to Program Structure

Metrics are assigned to program instructions using the call stack that is recorded with the event-specific data. If the information is available, each instruction is mapped to a line of source code and the metrics assigned to that instruction are also assigned to the line of source code. See [Chapter 5](#) for a more detailed explanation of how this is done.

In addition to source code and instructions, metrics are assigned to higher level objects: functions and load objects. The call stack contains information on the sequence of function calls made to arrive at the instruction address recorded when a profile was taken. The Performance Analyzer uses the call stack to compute metrics for each function in the program. These metrics are called function-level metrics.

## Function-Level Metrics: Exclusive, Inclusive, and Attributed

The Performance Analyzer computes three types of function-level metrics: exclusive metrics, inclusive metrics and attributed metrics.

- Exclusive metrics for a function are calculated from events which occur inside the function itself: they exclude metrics coming from calls to other functions.
- Inclusive metrics are calculated from events which occur inside the function and any functions it calls: they include metrics coming from calls to other functions.
- Attributed metrics tell you how much of an inclusive metric came from calls from or to another function: they attribute metrics to another function.

For a function at the bottom of a particular call stack (the leaf function), the exclusive and inclusive metrics are the same, because the function makes no calls to other functions.

Exclusive and inclusive metrics are also computed for load objects. Exclusive metrics for a load object are calculated by summing the function-level metrics over all functions in the load object. Inclusive metrics for load objects are calculated in the same way as for functions.

Exclusive and inclusive metrics for a function give information about all recorded paths through the function. Attributed metrics give information about particular paths through a function. They show how much of a metric came from a particular function call. The two functions involved in the call are described as a *caller* and a *callee*. For each function in the call tree:

- The attributed metrics for a function's callers tell you how much of the function's inclusive metric was due to calls from each caller. The attributed metrics for the callers sum to the function's inclusive metric.
- The attributed metrics for a function's callees tell you how much of the function's inclusive metric came from calls to each callee. Their sum plus the function's exclusive metric equals the function's inclusive metric.

Comparison of attributed and inclusive metrics for the caller or the callee gives further information:

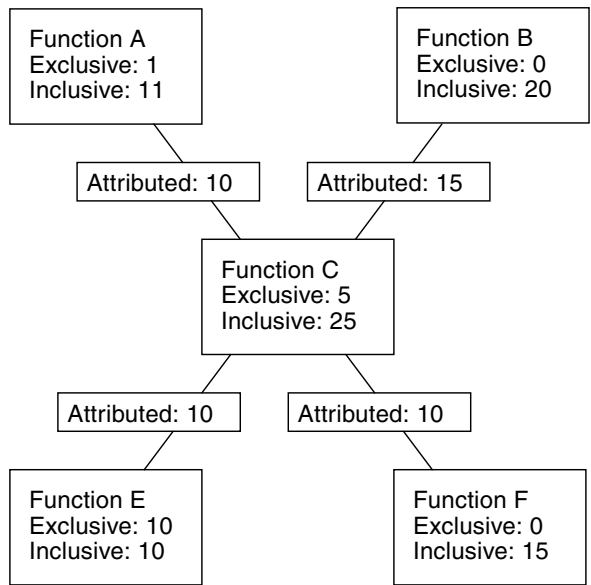
- The difference between a caller's attributed metric and its inclusive metric tells you how much of the metric came from calls to other functions and from work in the caller itself.
- The difference between a callee's attributed metric and its inclusive metric tells you how much of the callee's inclusive metric came from calls to it from other functions.

To locate places where you could improve the performance of your program:

- Use exclusive metrics to locate functions that have high metric values.
- Use inclusive metrics to determine which call sequence in your program was responsible for high metric values.
- Use attributed metrics to trace a particular call sequence to the function or functions that are responsible for high metric values.

## Interpreting Attributed Metrics: An Example

Exclusive, inclusive and attributed metrics are illustrated in [FIGURE 2-1](#), which contains a fragment of a call tree. The focus is on the central function, function C. There may be calls to other functions which do not appear in this figure.



**FIGURE 2-1** Call Tree Illustrating Exclusive, Inclusive, and Attributed Metrics

Function C calls two functions, function E and function F, and attributes 10 units of its inclusive metric to function E and 10 units to function F. These are the callee attributed metrics. Their sum (10+10) added to the exclusive metric of function C (5) equals the inclusive metric of function C (25).

The callee attributed metric and the callee inclusive metric are the same for function E but different for function F. This means that function E is only called by function C but function F is called by some other function or functions. The exclusive metric and the inclusive metric are the same for function E but different for function F. This means that function F calls other functions, but function E does not.

Function C is called by two functions: function A and function B, and attributes 10 units of its inclusive metric to function A and 15 units to function B. These are the caller attributed metrics. Their sum (10+15) equals the inclusive metric of function C.

The caller attributed metric is equal to the difference between the inclusive and exclusive metric for function A, but it is not equal to this difference for function B. This means that function A only calls function C, but function B calls other functions besides function C. (In fact, function A might call other functions but the time is so small that it does not appear in the experiment.)



## How Recursion Affects Function-Level Metrics

Recursive function calls, whether direct or indirect, complicate the calculation of metrics. The Performance Analyzer displays metrics for a function as a whole, not for each invocation of a function: the metrics for a series of recursive calls must therefore be compressed into a single metric. This does not affect exclusive metrics, which are calculated from the function at the bottom of the call stack (the leaf function), but it does affect inclusive and attributed metrics.

Inclusive metrics are computed by adding the exclusive metric for the leaf function to the inclusive metric of the functions in the call stack. To ensure that the metric is not counted multiple times in a recursive call stack, the exclusive metric for the leaf function is only added to the inclusive metric for each unique function.

Attributed metrics are computed from inclusive metrics. In the simplest case of recursion, a recursive function has two callers: itself and another function (the initiating function). If all the work is done in the final call, the inclusive metric for the recursive function will be attributed to itself and not to the initiating function. This is because the inclusive metric for all the higher invocations of the recursive function are regarded as zero to avoid multiple counting of the metric. The initiating function, however, correctly attributes to the recursive function as a callee the portion of its inclusive metric due to the recursive call.



## Collecting Performance Data

---

The first stage of performance analysis is data collection. This chapter describes what is required for data collection, where the data is stored, how to collect data and how to manage the data collection. For more information about the data itself, see [Chapter 2](#).

This chapter covers the following topics.

- [Compiling and Linking Your Program](#)
- [Preparing Your Program for Data Collection and Analysis](#)
- [Limitations on Data Collection](#)
- [Where the Data Is Stored](#)
- [Estimating Storage Requirements](#)
- [Collecting Data](#)
- [Collecting Data Using the `collect` Command](#)
- [Collecting Data Using the `dbx collector` Subcommands](#)
- [Collecting Data From a Running Process](#)
- [Collecting Data From MPI Programs](#)
- [Using `collect` With `ppgsz`](#)

---

## Compiling and Linking Your Program

You can collect and analyze data for a program compiled with almost any option, but some choices affect what you can collect or what you can see in the Performance Analyzer. The issues that you should take into account when you compile and link your program are described in the following subsections.

## Source Code Information

To see source code in annotated Source and Disassembly, and source lines in the Lines analyses, you must compile the source files of interest with the `-g` compiler option (`-g0` for C++ to enable front-end inlining) to generate debug symbol information. The format of the debug symbol information can be either STABS or DWARF2, as specified by `-xdebugformat=(stabs|dwarf)`.

To prepare compilation objects with debug information which allows dataspaces hardware counter profiles, currently only for the C compiler for SPARC<sup>®</sup>, compile by specifying `-xhwcprof -xdebugformat=dwarf` and any level of optimization. (Currently, this functionality is not available without optimization.) To see program data objects in Data Objects analyses, also add `-g` to obtain full symbolic information.

Executables and libraries built with DWARF format debugging symbols automatically include a copy of each constituent object file's debugging symbols. This is also true for STABS format debugging symbols if they are linked with the `-xs` option which leaves STABS symbols in the various object files as well as the executable. This is particularly useful if you need to move or remove the object files. With all of the debugging symbols in the executables and libraries themselves, it is easier to move the experiment and the program-related files to a new location.

## Static Linking

When you compile your program, you must not disable dynamic linking, which is done with the `-dn` and `-Bstatic` compiler options. If you try to collect data for a program that is entirely statically linked, the Collector prints an error message and does not collect data. This is because the collector library, among others, is dynamically loaded when you run the Collector.

You should not statically link any of the system libraries. If you do, you might not be able to collect any kind of tracing data. Nor should you link to the Collector library, `libcollector.so`.

## Optimization

If you compile your program with optimization turned on at some level, the compiler can rearrange the order of execution so that it does not strictly follow the sequence of lines in your program. The Performance Analyzer can analyze experiments collected on optimized code, but the data it presents at the disassembly level is often difficult to relate to the original source code lines. In addition, the call

sequence can appear to be different from what you expect if the compiler performs tail-call optimizations. Optimization may cause unwind failures. See [“Tail-Call Optimization” on page 120](#) for more information.

## Compiling Java Programs

No special action is required for compiling Java programs with `javac`.

---

# Preparing Your Program for Data Collection and Analysis

For most programs, you do not need to do anything special to prepare your program for data collection and analysis. You should read one or more of the subsections below if your program does any of the following:

- Installs a signal handler
- Explicitly dynamically loads a system library
- Dynamically compiles functions
- Creates descendant processes
- Uses the asynchronous I/O library
- Uses the profiling timer or hardware counter API directly
- Calls `setuid(2)` or executes a `setuid` file.

Also, if you want to control data collection from your program you should read the relevant subsection.

## Using Dynamically Allocated Memory

Many programs rely on dynamically-allocated memory, using features such as:

- `malloc`, `valloc`, `alloca` (C/C++)
- `new` (C++)
- stack local variables (Fortran)
- `MALLOC`, `MALLOC64` (Fortran)

You must take care to ensure that a program does not rely on the initial contents of dynamically allocated memory, unless the memory allocation method is explicitly documented as setting an initial value: for example, compare the descriptions of `calloc` and `malloc` in the man page for `malloc(3C)`.

Occasionally, a program that uses dynamically-allocated memory may appear to work correctly when run alone, but may fail when run with performance data collection enabled. Symptoms might include unexpected floating point behavior, segmentation faults, or application-specific error messages.

Such behavior may occur if the uninitialized memory is, by chance, set to a benign value when the application is run alone, but is set to a different value when the application is run in conjunction with the performance data collection tools. In such cases, the performance tools are not at fault. Any application that relies on the contents of dynamically allocated memory has a latent bug: an operating system is at liberty to provide any content whatsoever in dynamically allocated memory, unless explicitly documented otherwise. Even if an operating system happens to always set dynamically allocated memory to a certain value today, such latent bugs might cause unexpected behavior with a later revision of the operating system, or if the program is ported to a different operating system in the future.

here are some tools that may help in finding such latent bugs:

- `f95 -xcheck=init_local`

For more information, see the *Fortran User's Guide* or the `f95(1)` man page

- `lint`

For more information, see the *C User's Guide* or the `lint(1)` man page

- Runtime checking under `dbx`

For more information, see the *Debugging a Program With dbx* manual or the `dbx(1)` man page.

- Purify

## Using System Libraries

The Collector interposes on functions from various system libraries, to collect tracing data and to ensure the integrity of data collection. The following list describes situations in which the Collector interposes on calls to library functions.

- Collection of synchronization wait tracing data. The Collector interposes on functions from the threads library, `libthread.so`.
- Collection of heap tracing data. The Collector interposes on the functions `malloc`, `realloc`, `memalign` and `free`. Versions of these functions are found in the C standard library, `libc.so` and also in other libraries such as `libmalloc.so` and `libmtmalloc.so`.
- Collection of MPI tracing data. The Collector interposes on functions from the MPI library, `libmpi.so`.
- Ensuring the integrity of clock data. The Collector interposes on `setitimer` and prevents the program from using the profiling timer.

- Ensuring the integrity of hardware counter data. The Collector interposes on functions from the hardware counter library, `libcpc.so` and prevents the program from using the counters. Calls from the program to functions from this library return with a return value of `-1`.
- Enabling data collection on descendant processes. The Collector interposes on the functions `fork(2)`, `fork1(2)`, `vfork(2)`, `fork(3F)`, `system(3C)`, `system(3F)`, `sh(3F)`, `popen(3C)`, and `exec(2)` and its variants. Calls to `vfork` are replaced internally by calls to `fork1`. These interpositions are only done for the `collect` command.
- Guaranteeing the handling of the `SIGPROF` and `SIGEMT` signals by the Collector. The Collector interposes on `sigaction` to ensure that its signal handler is the primary signal handler for these signals.

There are some circumstances in which the interposition does not succeed:

- Statically linking a program with any of the libraries that contain functions that are interposed.
- Attaching `dbx` to a running application that does not have the collector library preloaded.
- Dynamically loading one of these libraries and resolving the symbols by searching only within the library.

The failure of interposition by the Collector can cause loss or invalidation of performance data.

## Using Signal Handlers

The Collector uses two signals to collect profiling data, `SIGPROF` and `SIGEMT`. The Collector installs a signal handler for each of these signals, which intercept and process the signals, but pass on signals they do not use to any other signal handlers that are installed. If a program installs its own signal handler for these signals, the Collector re-installs its signal handler as the primary handler to guarantee the integrity of the performance data.

The `collect` command can also use user-specified signals for pausing and resuming data collection and for recording samples. These signals are not protected by the Collector although a warning is written to the experiment if a user handler is installed. It is the responsibility of the user to ensure that there is no conflict between use of the specified signals by the Collector and any use made by the application of the same signals.

The signal handlers installed by the Collector set a flag that ensures that system calls are not interrupted for signal delivery. This flag setting could change the behavior of the program if the program's signal handler sets the flag to permit interruption of system calls. One important example of a change in behavior occurs for the

asynchronous I/O library, `libaio.so`, which uses `SIGPROF` for asynchronous cancel operations, and which does interrupt system calls. If the collector library, `libcollector.so`, is installed, the cancel signal arrives late.

If you attach `dbx` to a process without preloading the collector library and enable performance data collection, and the program subsequently installs its own signal handler, the Collector does not re-install its own signal handler. In this case, the program's signal handler must ensure that the `SIGPROF` and `SIGEMT` signals are passed on so that performance data is not lost. If the program's signal handler interrupts system calls, both the program behavior and the profiling behavior will be different from when the collector library is preloaded.

## Using `setuid`

There are restrictions enforced by the dynamic loader that make it difficult to use `setuid(2)` and collect performance data. If your program calls `setuid` or executes a `setuid` file, it is likely that the Collector cannot write an experiment file because it lacks the necessary permissions for the new user ID.

## Program Control of Data Collection

If you want to control data collection from your program, the Collector shared library, `libcollector.so` contains some API functions that you can use. The functions are written in C. A Fortran interface is also provided. Both C and Fortran interfaces are defined in header files that are provided with the library.

The API functions are defined as follows.

```
void collector_sample(char *name);
void collector_pause(void);
void collector_resume(void);
void collector_thread_pause(unsigned int t);
void collector_thread_resume(unsigned int t);
void collector_terminate_expt(void);
```

Similar functionality is provided for Java programs by the `CollectorAPI` class, which is described in [“The Java Interface” on page 50](#).

## The C and C++ Interface

There are two ways to access the C and C++ interface.



- One way is to include `collectorAPI.h` and link with `-lcollectorAPI` (which contains real functions to check for the existence of the underlying `libcollector.so` API functions).

This way requires that you link with an API library, and works under all circumstances. If no experiment is active, the API calls are ignored.

- The second way is to include `libcollector.h` (which contains macros that check for the existence of the underlying `libcollector.so` API functions.)

This way works when used in the main executable, and when data collection is started at the same time the program starts. This way does not always work when `dbx` is used to attach to the process, nor when used from within a shared library that is `dlopen'd` by the process. This second way is provided for backward compatibility.

---

**Caution** – Do not link a program in any language with `-lcollector`. If you do, the Collector can exhibit unpredictable behavior.

---

## The Fortran Interface

The Fortran API `libfcollector.h` file defines the Fortran interface to the library. The application must be linked with `-lcollectorAPI` to use this library. (An alternate name for the library, `-lfcollector`, is provided for backward compatibility.) The Fortran API provides the same features as the C and C++ API, excluding the dynamic function and thread pause and resume calls.

Insert the following statement to use the API functions for Fortran:

```
include "libfcollector.h"
```

---

**Caution** – Do not link a program in any language with `-lcollector`. If you do, the Collector can exhibit unpredictable behavior.

---

## The Java Interface

Use the following statement to import the CollectorAPI class and access the Java API. Note however that your application must be invoked with a classpath pointing to <installation-directory>/lib/collector.jar where <installation-directory> is the directory into which the Sun compilers and tools are installed.

```
import com.sun.forte.st.collector.CollectorAPI;
```

The Java CollectorAPI methods are defined as follows:

```
CollectorAPI.sample(String name)
CollectorAPI.pause()
CollectorAPI.resume()
CollectorAPI.threadPause(Thread thread)
CollectorAPI.threadResume(Thread thread)
CollectorAPI.terminate()
```

The Java API includes the same functions as the Fortran API, excluding the dynamic function API.

The C include-file `libcollector.h` contains macros that bypass the calls to the real API functions if data is not being collected. In this case the functions are not dynamically loaded. However, using these macros is risky because the macros do not work well under some circumstances. It is safer to use `collectorAPI.h` because it does not use macros. Rather, it refers directly to the functions.

The Fortran API subroutines call the C API functions if performance data is being collected, otherwise they return. The overhead for the checking is very small and should not significantly affect program performance.

To collect performance data you must run your program using the Collector, as described later in this chapter. Inserting calls to the API functions does not enable data collection.

If you intend to use the API functions in a multithreaded program, you should ensure that they are only called by one thread. With the exception of `collector_thread_pause()` and `collector_thread_resume()`, the API functions perform actions that apply to the process and not to individual threads. If each thread calls the API functions, the data that is recorded might not be what you expect. For example, if `collector_pause()` or `collector_terminate_expt()` is called by one thread before the other threads have reached the same point in the program, collection is paused or terminated for all threads, and data can be lost from the threads that were executing code before the API call. To control data collection at the level of the individual threads, use the `collector_thread_pause()` and

`collector_thread_resume()` functions. There are two ways of using these functions: by having one master thread make all the calls for all threads, including itself; or by having each thread make calls only for itself. Any other usage can lead to unpredictable results.

## The C, C++, Fortran, and Java API Functions

The descriptions of the API functions follow.

- **C and C++:** `collector_sample(char *name)`

**Fortran:** `collector_sample(string)`

**Java:** `CollectorAPI.sample(String)`

Record a sample packet and label the sample with the given name or string. The label is displayed by the Performance Analyzer in the Event tab. The Fortran argument `string` is of type character.

Sample points contain data for the process and not for individual threads. In a multithreaded application, the `collector_sample()` API function ensures that only one sample is written if another call is made while it is recording a sample. The number of samples recorded can be less than the number of threads making the call.

The Performance Analyzer does not distinguish between samples recorded by different mechanisms. If you want to see only the samples recorded by API calls, you should turn off all other sampling modes when you record performance data.

- **C, C++, Fortran:** `collector_pause()`

**Java:** `CollectorAPI.pause()`

Stop writing event-specific data to the experiment. The experiment remains open, and global data continues to be written. The call is ignored if no experiment is active or if data recording is already stopped. This function stops the writing of all event-specific data even if it is enabled for specific threads by the `collector_thread_resume()` function.

- **C, C++, Fortran:** `collector_resume()`

**Java:** `CollectorAPI.resume()`

Resume writing event-specific data to the experiment after a call to `collector_pause()`. The call is ignored if no experiment is active or if data recording is active.

- **C and C++ only:** `collector_thread_pause(unsigned int t)`

**Java:** `CollectorAPI.threadPause(Thread)`

Stop writing event-specific data from the thread specified in the argument list to the experiment. The argument `t` is the POSIX thread identifier. If the experiment is already terminated, or no experiment is active, or writing of data for that thread is already turned off, the call is ignored. This function stops the writing of data from the specified thread even if the writing of data is globally enabled. By default, recording of data for individual threads is turned on.

- **C and C++ only:** `collector_thread_resume(unsigned int t)`

**Java:** `CollectorAPI.threadResume(Thread)`

Resume writing event-specific data from the thread specified in the argument list to the experiment. The argument `t` is the POSIX thread identifier. If the experiment is already terminated, or no experiment is active, or writing of data for that thread is already turned on, the call is ignored. Data is written to the experiment only if the writing of data is globally enabled as well as enabled for the thread.

- **C, C++, Fortran:** `collector_terminate_expt()`

**Java:** `CollectorAPI.terminate`

Terminate the experiment whose data is being collected. No further data is collected, but the program continues to run normally. The call is ignored if no experiment is active.

## Dynamic Functions and Modules

If your C program or C++ program dynamically compiles functions into the data space of the program, you must supply information to the Collector if you want to see data for the dynamic function or module in the Performance Analyzer. The information is passed by calls to collector API functions. The definitions of the API functions are as follows.

```
void collector_func_load(char *name, char *alias,
    char *sourcename, void *vaddr, int size, int lntsize,
    Lineno *lntable);
void collector_func_unload(void *vaddr);
```

You do not need to use these API functions for Java™ methods that are compiled by the Java HotSpot™ virtual machine, for which a different interface is used. The Java interface provides the name of the method that was compiled to the Collector. You can see function data and annotated disassembly listings for Java compiled methods, but not annotated source listings.

The descriptions of the API functions follow.

- `collector_func_load()`

Pass information about dynamically compiled functions to the Collector for recording in the experiment. The parameter list is described in the following table.

**TABLE 3-1** Parameter List for `collector_func_load()`

Parameter	Definition
<code>name</code>	The name of the dynamically compiled function that is used by the performance tools. The name does not have to be the actual name of the function. The name need not follow any of the normal naming conventions of functions, although it should not contain embedded blanks or embedded quote characters.
<code>alias</code>	An arbitrary string used to describe the function. It can be <code>NULL</code> . It is not interpreted in any way, and can contain embedded blanks. It is displayed in the Summary tab of the Analyzer. It can be used to indicate what the function is, or why the function was dynamically constructed.
<code>sourcename</code>	The path to the source file from which the function was constructed. It can be <code>NULL</code> . The source file is used for annotated source listings.
<code>vaddr</code>	The address at which the function was loaded.
<code>size</code>	The size of the function in bytes.
<code>lntsize</code>	A count of the number of entries in the line number table. It should be zero if line number information is not provided.
<code>lntable</code>	A table containing <code>lntsize</code> entries, each of which is a pair of integers. The first integer is an offset, and the second entry is a line number. All instructions between an offset in one entry and the offset given in the next entry are attributed to the line number given in the first entry. Offsets must be in increasing numeric order, but the order of line numbers is arbitrary. If <code>lntable</code> is <code>NULL</code> , no source listings of the function are possible, although disassembly listings are available.

■ `collector_func_unload()`

Inform the collector that the dynamic function at the address `vaddr` has been unloaded.

---

## Limitations on Data Collection

This section describes the limitations on data collection that are imposed by the hardware, the operating environment, the way you run your program or by the Collector itself.

There are no limitations on simultaneous collection of different data types: you can collect any data type with any other data type.

## Limitations on Clock-Based Profiling

The minimum value of the profiling interval and the resolution of the clock used for profiling depend on the particular operating environment. The maximum value is set to 1 second. The value of the profiling interval is rounded down to the nearest multiple of the clock resolution. The minimum and maximum value and the clock resolution can be found by typing the `collect` command with no arguments.

The system clock is used for profiling in the Solaris 7 Operating System and versions of the Solaris 8 Operating System. It has a resolution of 10 milliseconds, unless you choose to enable the high-resolution system clock. If you have root privilege, you can do this by adding the following line to the file `/etc/system`, and then rebooting.

```
set hires_tick=1
```

In the Solaris 9 Operating System and later versions of the Solaris 8 Operating System, it is not necessary to enable the high-resolution system clock for high-resolution profiling.

## Runtime Distortion and Dilation with Clock-profiling

Clock profiling records data when a SIGPROF signal is delivered to the target. It will cause dilation to process that signal, and unwind the callstack. The deeper the callstack, and the more frequent the signals, the greater the dilation. To a limited extent, clock-profiling will show some distortion, deriving from greater dilation for those parts of the program executing with the deepest stacks.

## Limitations on Collection of Tracing Data

You cannot collect any kind of tracing data from a program that is already running unless the Collector library, `libcollector.so`, has been preloaded. See [“Collecting Data From a Running Process” on page 76](#) for more information.

## Runtime Distortion and Dilation with Tracing

Tracing data will dilate the run, in proportion to the number of events that are traced. If done with clock-profiling, the clock data will be distorted by the dilation induced by tracing events.

## Limitations on Hardware-Counter Overflow Profiling

There are several limitations on hardware counter overflow profiling:

- You can only collect hardware-counter overflow data on processors that have hardware counters and that support overflow profiling. On other systems, hardware-counter overflow profiling is disabled. UltraSPARC® processors prior to the UltraSPARC® III processor family do not support hardware-counter overflow profiling.
- You cannot collect hardware-counter overflow data with versions of the operating environment that precede the release of the Solaris 8 Operating System.
- You can record data for at most two hardware counters in an experiment. To record data for more than two hardware counters or for counters that use the same register you must run separate experiments.
- You cannot collect hardware-counter overflow data on a system while `cpustat(1)` is running, because `cpustat` takes control of the counters and does not let a user process use the counters. If `cpustat` is started during data collection, the hardware counter overflow profiling is terminated.
- You cannot use the hardware counters in your own code with the `libcpc(3)` API if you are doing hardware-counter overflow profiling. The Collector interposes on the `libcpc` library functions and returns with a return value of `-1` if the call did not come from the Collector.
- If you try to collect hardware counter data on a running program that is using the hardware counter library by attaching `dbx` to the process, the experiment may be corrupted.

---

**Note** – To view a list of all available counters, run `collect` with no arguments.

---

# Runtime Distortion and Dilation With Hardware-Counter Overflow Profiling

Hardware-counter profiling records data when a SIGEMT is delivered to the target. It will cause dilation to process that signal, and unwind the callstack. Unlike clock profiling, for some hardware counters, different parts of the program may generate events more rapidly than other parts, and will show dilation in that part of the code. Any part of the program that generates such events very rapidly may be significantly distorted. Similarly, some events may be generated in one thread disproportionately to the other threads.

## Limitations on Data Collection for Descendant Processes

You can collect data on descendant processes subject to the following limitations:

- If you want to collect data for all descendant processes that are followed by the Collector, you must use the `collect` command with the `-F on` option.
- You can collect data automatically for calls to `fork` and its variants and `exec` and its variants. Calls to `system`, `popen`, and `sh` are not followed by the Collector.
- If you want to collect data for individual descendant processes, you must attach `dbx` to the process. See [Appendix “Collecting Data From a Running Process” on page 76](#) for more information.
- If you want to collect data for individual descendant processes, or those created by `system`, `popen`, `sh`, etc., you must use a separate `dbx` to attach to each process and enable the collector.

## Limitations on Java Profiling

You can collect data on Java programs subject to the following limitations:

- You should use a version of the Java™ 2 Software Development Kit no earlier than 1.4.2\_02. Do not use versions 1.4.2 or 1.4.2\_01 which may crash. The path to the Java virtual machine<sup>1</sup> should be specified in one of the following four environment variables: `JDK_1_4_HOME`, `JDK_HOME`, `JAVA_PATH`, `PATH`. The Collector verifies that the version of `java` it finds in these environment variables is an ELF executable, and if it is not, an error message is printed, indicating which environment variable was used, and the full path name that was tried.

---

1. The terms “Java virtual machine” and “JVM” mean a virtual machine for the Java platform.



- You must use the `collect` command to collect data. You cannot use the `dbx` collector subcommands or the data collection capabilities of the IDE.
- If you want to use the 64 bit JVM™, it must either be the default, or you must specify the path to it when you collect data. Do not use `java -d64` to collect data using the 64 bit JVM. If you do, no data is collected.

Using JVM versions earlier than 1.4.2\_02 compromises the data as follows:

- **JVM 1.4.2\_01:** This version of the JVM may crash during data collection.
- **JVM 1.4.2\_02:** This version of the JVM may crash during data collection.
- **JVM 1.4.1:** The Java representation is correctly recorded and shown, but all JVM housekeeping is shown as the JVM functions themselves. Some of the time spent executing JVM code in data space is shown with names for the code regions as supplied by the JVM. A significant amount of time will be shown in the `<Unknown>` function, since some of the code regions created by the JVM are not named. In addition, there are various bugs in JVM 1.4.1 that may cause the program being profiled to crash.
- **JVM 1.4.0:** No Java representation is possible, and a significant amount of time is shown in `<Unknown>`. HotSpot-compiled functions are shown by name in the machine representation.
- **JVMs earlier than 1.4.0:** Profiling Java applications with JVMs earlier than 1.4.0 is not supported.

## Runtime Performance Distortion and Dilation for Applications Written in the Java Programming Language

Java profiling uses the JVMPI interface, which can cause some distortion and dilation of the run. For clock- and hwc-profiling, the data collection process makes various calls into the JVM, and handles profiling events in signal handlers. The overhead of these routines, and the cost of writing the experiments to disk will dilate the runtime of the Java program. Such dilation is estimated to be less than 10%.

In addition, although the default garbage collector supports JVMPI, there are other garbage collectors that do not. Any data-collection run specifying such a garbage collector will get a fatal error.

For Heap profiling, the data collection process uses JVMPI events describing memory allocation and garbage collection, which can cause significant dilation in runtime. Most Java applications generate many of these events, which will lead to large experiments, and scalability problems processing the data. Furthermore, if these events are requested, the garbage collector disables some inlined allocations, costing additional CPU time for the longer allocation path.

For synchronization tracing, data collection uses other JVMPI events, which will cause dilation in proportion to the amount of monitor contention in the application.

---

## Where the Data Is Stored

The data collected during one execution of your application is called an experiment. The experiment consists of a set of files that are stored in a directory. The name of the experiment is the name of the directory.

In addition to recording the experiment data, the Collector creates its own archives of the load objects used by the program. These archives contain the addresses, sizes and names of each object file and each function in the load object, as well as the address of the load object and a time stamp for its last modification.

Experiments are stored by default in the current directory. If this directory is on a networked file system, storing the data takes longer than on a local file system, and can distort the performance data. You should always try to record experiments on a local file system if possible. You can change the storage location when you run the Collector.

Experiments for descendant processes are stored inside the experiment for the founder process.

## Experiment Names

The default name for a new experiment is `test.1.er`. The suffix `.er` is mandatory: if you give a name that does not have it, an error message is displayed and the name is not accepted.

If you choose a name with the format *experiment.n.er*, where *n* is a positive integer, the Collector automatically increments *n* by one in the names of subsequent experiments—for example, `mytest.1.er` is followed by `mytest.2.er`, `mytest.3.er`, and so on. The Collector also increments *n* if the experiment already exists, and continues to increment *n* until it finds an experiment name that is not in use. If the experiment name does not contain *n* and the experiment exists, the Collector prints an error message.

Experiments can be collected into groups. The group is defined in an experiment group file, which is stored by default in the current directory. The experiment group file is a plain text file with a special header line and an experiment name on each subsequent line. The default name for an experiment group file is `test.erg`. If the

name does not end in `.erg`, an error is displayed and the name is not accepted. Once you have created an experiment group, any experiments you run with that group name are added to the group.

You can create an experiment group file by creating a plain text file whose first line is

```
#analyzer experiment group
```

and adding the names of the experiments on subsequent lines. The name of the file must end in `.erg`.

The default experiment name is different for experiments collected from MPI programs, which create one experiment for each MPI process. The default experiment name is `test.m.er`, where `m` is the MPI rank of the process. If you specify an experiment group `group.erg`, the default experiment name is `group.m.er`. If you specify an experiment name, it overrides these defaults. See [“Collecting Data From MPI Programs” on page 78](#) for more information.

Experiments for descendant processes are named with their lineage as follows. To form the experiment name for a descendant process, an underscore, a code letter and a number are added to the stem of its creator’s experiment name. The code letter is `f` for a fork and `x` for an exec. The number is the index of the fork or exec (whether successful or not). For example, if the experiment name for the founder process is `test.1.er`, the experiment for the child process created by the third call to `fork` is `test.1.er/_f3.er`. If that child process calls `exec` successfully, the experiment name for the new descendant process is `test.1.er/_f3_x1.er`.

## Moving Experiments

If you want to move an experiment to another computer to analyze it, you should be aware of the dependencies of the analysis on the operating environment in which the experiment was recorded.

The archive files contain all the information necessary to compute metrics at the function level and to display the timeline. However, if you want to see annotated source code or annotated disassembly code, you must have access to versions of the load objects or source files that are identical to the ones used when the experiment was recorded.

The Performance Analyzer searches for the source, object and executable files in the following locations in turn, and stops when it finds a file of the correct basename:

- The archive directories of experiments.
- The current working directory.
- The absolute pathname as recorded in the executables or compilation objects.

To ensure that you see the correct annotated source code and annotated disassembly code for your program, you can copy the source code, the object files and the executable into the experiment before you move or copy the experiment. If you don't want to copy the object files, you can link your program with `-xs` to ensure that the information on source lines and file locations are inserted into the executable. You can automatically copy the load objects into the experiment using the `-A` option of the `collect` command or the `dbx collector archive` command.

---

## Estimating Storage Requirements

In this section some guidelines are given for estimating the amount of disk space needed to record an experiment. The size of the experiment depends directly on the size of the data packets and the rate at which they are recorded, the number of LWPs used by the program, and the execution time of the program.

The data packets contain event-specific data and data that depends on the program structure (the call stack). The amount of data that depends on the data type is approximately 50 to 100 bytes. The call stack data consists of return addresses for each call, and contains 4 bytes (8 bytes on 64 bit SPARC architecture) per address. Data packets are recorded for each LWP in the experiment. Note that for Java programs, there will be two callstacks of interest: the Java callstack and the machine callstack, which therefore will result in more data being written to disk.

The rate at which profiling data packets are recorded is controlled by the profiling interval for clock data and by the overflow value for hardware counter data. However, the choice of these parameters also affects the data quality and the distortion of program performance due to the data collection overhead. Smaller values of these parameters give better statistics but also increase the overhead. The default values of the profiling interval and the overflow value have been carefully chosen as a compromise between obtaining good statistics and minimizing the overhead. Smaller values also mean more data.

For a clock-based profiling experiment with a profiling interval of 10ms and a small call stack, such that the packet size is 100 bytes, data is recorded at a rate of 10 kbytes/sec per LWP. For a hardware counter overflow profiling experiment collecting data for CPU cycles and instructions executed on a 750MHz processor with an overflow value of 1000000 and a packet size of 100 bytes, data is recorded at a rate of 150 kbytes/sec per LWP. Applications that have call stacks with a depth of hundreds of calls could easily record data at ten times these rates.

Your estimate of the size of the experiment should also take into account the disk space used by the archive files, which is usually a small fraction of the total disk space requirement (see the previous section). If you are not sure how much space you need, try running your experiment for a short time. From this test you can

obtain the size of the archive files, which are independent of the data collection time, and scale the size of the profile files to obtain an estimate of the size for the full-length experiment.

As well as allocating disk space, the Collector allocates buffers in memory to store the profile data before writing it to disk. There is currently no way to specify the size of these buffers. If the Collector runs out of memory, you should try to reduce the amount of data collected.

If your estimate of the space required to store the experiment is larger than the space you have available, you can consider collecting data for part of the run rather than the whole run. You can do this with the `collect` command, with the `dbx collector` subcommands, or by inserting calls in your program to the collector API. You can also limit the total amount of profiling and tracing data collected with the `collect` command or with the `dbx collector` subcommands.

---

**Note** – The Performance Analyzer cannot read more than 2 GB of performance data.

---

## Collecting Data

You can collect performance data in either the standalone Performance Analyzer or the Analyzer module in several ways:

- Using the `collect` command from the command line (see [“Collecting Data Using the `collect` Command” on page 62](#) and the `collect(1)` man page). The `collect` command-line tool has smaller data collection overheads than `dbx` or the Collector dialog in the Debugger in the IDE so this method can be superior to the others.
- Using the Performance Tools Collect dialog in the Performance Analyzer (see [“Collecting Performance Data Using the Performance Tools Collector”](#) in the Performance Analyzer online help)
- Using the Collector dialog in the Debugger (see [“Collecting Performance Data Using the Debugger”](#) in the Performance Analyzer online help)
- Using the `collector` command from the `dbx` command line (see [“Collecting Data Using the `dbx collector` Subcommands” on page 70](#) and [“collector Command”](#) in the Debugging online help in the IDE)

The following data collection capabilities are available only with the Performance Tools Collect dialog and the `collect` command:

- Collecting data on Java™ programs. If you try to collect data on a Java program with Collector Dialog in the Debugger in the IDE or with the `collector` command in `dbx`, the information that is collected is for the Java Virtual Machine, not the Java program.
- Collecting data automatically on descendant processes.

---

## Collecting Data Using the `collect` Command

To run the Collector from the command line using the `collect` command, type the following.

```
% collect collect-options program program-arguments
```

Here, *collect-options* are the `collect` command options, *program* is the name of the program you want to collect data on, and *program-arguments* are its arguments.

If no command arguments are given, the default is to turn on clock-based profiling with a profiling interval of 10 milliseconds.

To obtain a list of options and a list of the names of any hardware counters that are available for profiling, type the `collect` command with no arguments.

```
% collect
```

For a description of the list of hardware counters, see [“Hardware-Counter Overflow Profiling Data” on page 30](#). See also [“Limitations on Hardware-Counter Overflow Profiling” on page 55](#).

## Data Collection Options

These options control the types of data that are collected. See [“What Data the Collector Collects” on page 28](#) for a description of the data types.

If no data collection options are given, the default is `-p on`, which enables clock-based profiling with the default profiling interval of 10 milliseconds. The default is turned off by the `-h` option but not by any of the other data collection options.

If clock-based profiling is explicitly disabled, and neither any kind of tracing nor hardware counter overflow profiling is enabled, the `collect` command prints a warning message, and collects global data only.

### `-p option`

Collect clock-based profiling data. The allowed values of *option* are:

- `off` – Turn off clock-based profiling.
- `on` – Turn on clock-based profiling with the default profiling interval of 10 milliseconds.
- `lo[w]` – Turn on clock-based profiling with the low-resolution profiling interval of 100 milliseconds.
- `hi[gh]` – Turn on clock-based profiling with the high-resolution profiling interval of 1 millisecond. In the Solaris 7 Operating System and earlier versions of the Solaris 8 operating environment, high-resolution profiling must be explicitly enabled. See [“Limitations on Clock-Based Profiling” on page 54](#) for information on enabling high-resolution profiling.
- *value* – Turn on clock-based profiling and set the profiling interval to *value*. The default units for *value* are milliseconds. You can specify *value* as an integer or a floating-point number. The numeric value can optionally be followed by the suffix `m` to select millisecond units or `u` to select microsecond units. The value should be a multiple of the clock resolution. If it is larger but not a multiple it is rounded down. If it is smaller, a warning message is printed and it is set to the clock resolution.

Collecting clock-based profiling data is the default action of the `collect` command.

### `-h counter [ , value [ , counter2 [ , value2 ] ] ]`

Collect hardware counter overflow profiling data. The counter names *counter* and *counter2* can be one of the following:

- An aliased counter name
- An internal name, as used by `cputrack(1)`. If the counter can use either event register, the event register to be used can be specified by appending `/0` or `/1` to the internal name.

If two counters are specified, they must use different registers. If they do not use different registers, the `collect` command prints an error message and exits. Some counters can count on either register.

To obtain a list of available counters, type `collect` with no arguments in a terminal window. A description of the counter list is given in the section [“Hardware Counter Lists” on page 31](#).

If the hardware counter counts events that relate to memory access, the counter name can be prefixed with a + sign to turn on searching for the true PC of the instruction that caused the counter overflow. If the search is successful, the PC and effective address that was referenced are stored in the event data packet.

The overflow value is the number of events counted at which the hardware counter overflows and the overflow event is recorded. The overflow values can be specified using *value* and *value2*, which can be set to one of the following:

- *hi [gh]* – The high-resolution value for the chosen counter is used. The abbreviation *h* is also supported for compatibility with previous software releases.
- *lo [w]* – The low-resolution value for the chosen counter is used.
- *number* – The overflow value. Must be a positive integer.
- *on*, or a null string – The default overflow value is used.

The default is the normal threshold, which is predefined for each counter and which appears in the counter list. See also [“Limitations on Hardware-Counter Overflow Profiling” on page 55](#).

If you use the *-h* option without explicitly specifying a *-p* option, clock-based profiling is turned off. To collect both hardware counter data and clock-based data, you must specify both a *-h* option and a *-p* option.

### *-s option*

Collect synchronization wait tracing data. The allowed values of *option* are:

- *all* – Enable synchronization wait tracing with a zero threshold. This option forces all synchronization events to be recorded.
- *calibrate* – Enable synchronization wait tracing and set the threshold value by calibration at runtime. (Equivalent to *on*.)
- *off* – Disable synchronization wait tracing.
- *on* – Enable synchronization wait tracing with the default threshold, which is to set the value by calibration at runtime. (Equivalent to *calibrate*.)
- *value* – Set the threshold to *value*, given as a positive integer in microseconds.

Synchronization wait tracing data is not recorded for Java monitors.

### *-H option*

Collect heap tracing data. The allowed values of *option* are:

- *on* – Turn on tracing of heap allocation and deallocation requests.
- *off* – Turn off heap tracing.



Heap tracing is turned off by default.

### -m *option*

Collect MPI tracing data. The allowed values of *option* are:

- `on` – Turn on tracing of MPI calls.
- `off` – Turn off tracing of MPI calls.

MPI tracing is turned off by default.

See “[MPI Tracing Data](#)” on page 35 for more information about the MPI functions whose calls are traced and the metrics that are computed from the tracing data.

### -S *option*

Record sample packets periodically. The allowed values of *option* are:

- `off` – Turn off periodic sampling.
- `on` – Turn on periodic sampling with the default sampling interval of 1 second.
- *value* – Turn on periodic sampling and set the sampling interval to *value*. The interval value must be positive, and is given in seconds.

By default, periodic sampling at 1 second intervals is enabled.

## Experiment Control Options

### -F *option*

Control whether or not descendant processes should have their data recorded. The allowed values of *option* are:

- `on` – Record experiments on all descendant processes that are followed by the Collector.
- `off` – Do not record experiments on descendant processes.

The Collector follows processes created by calls to the functions `fork(2)`, `fork1(2)`, `fork(3F)`, `vfork(2)`, and `exec(2)` and its variants. The call to `vfork` is replaced internally by a call to `fork1`. The Collector does not follow processes created by calls to `system(3C)`, `system(3F)`, `sh(3F)`, and `popen(3C)`, and similar functions, and their associated descendant processes.

If you specify the `-F` on argument, the Collector opens a new experiment for each descendant process inside the founder experiment. These new experiments are named as follows.

An underscore, a letter, and a number are added to the experiment suffix in order to form the experiment name for a descendant process. The letter “*f*” indicates a fork and the letter “*x*” indicates an exec. The number is the index of the fork or exec (whether successful or not). For example, if the experiment name for the initial process is `test.1.er`, the experiment for the child process created by its third fork is `test.1.er/_f3.er`. If that child process execs a new image, the corresponding experiment name is `test.1.er/_f3_x1.er`.

The Analyzer and `er_print` automatically read experiments for descendant processes when the founder experiment is read, but the experiments for the descendant processes are not selected for data display.

To select the data for display from the command line, specify the path name, including the founder experiment name, and the descendant experiments name inside that directory, explicitly to either `er_print` or `analyzer`. For example, here’s what you specify to see the data for the third fork of the `test.1.er` experiment:

```
er_print test.1.er/_f3.er
```

```
analyzertest.1.er/_f3.er
```

Alternatively, you can prepare an experiment group file with the explicit names of the descendant experiments in which you are interested.

To select the data for display in the Analyzer, load `test.1.er` and select Filter data... from the View menu. You’ll see a list of experiments with only the founder experiment (`test.1.er`) checked. Uncheck it and check `_f3.er`.

## `-j` *option*

Enable Java profiling for a nonstandard Java installation, or choose whether to collect data on methods compiled by the Java HotSpot virtual machine. The allowed values of *option* are:

- `on` – Recognize methods compiled by the Java HotSpot virtual machine, and attempt to record Java stacks.
- `off` – Do not attempt to recognize methods compiled by the Java HotSpot virtual machine.

This option is not needed if you want to collect data on a `.class` file or a `.jar` file, provided that the path to the `java` executable is in one of the following environment variables: `JDK_1_4_HOME`, `JDK_HOME`, `JAVA_PATH`, or `PATH`. You can then specify *program* as the `.class` file or the `.jar` file, with or without the extension.

If you cannot define the path to `java` in any of these variables, or if you want to disable the recognition of methods compiled by the Java HotSpot virtual machine you can use this option. If you use this option, *program* must be a Java virtual machine whose version is not earlier than 1.4.2\_02. The `collect` command does not verify that *program* is a JVM machine, and collection can fail if it is not. However it does verify that *program* is an ELF executable, and if it is not, the `collect` command prints an error message.

If you want to collect data using the 64 bit JVM machine, you must not use the `-d64` option to `java` for a 32 bit JVM machine. If you do, no data is collected. Instead you must specify the path to the 64 bit JVM machine either in *program* or in one of the environment variables given in this section.

### `-l signal`

Record a sample packet when the signal named *signal* is delivered to the process.

The signal can be specified by the full signal name, by the signal name without the initial letters `SIG`, or by the signal number. Do not use a signal that is used by the program or that would terminate execution. Suggested signals are `SIGUSR1` and `SIGUSR2`. Signals can be delivered to a process by the `kill(1)` command.

If you use both the `-l` and the `-y` options, you must use different signals for each option.

If you use this option and your program has its own signal handler, you should make sure that the signal that you specify with `-l` is passed on to the Collector's signal handler, and is not intercepted or ignored.

See the `signal(3HEAD)` man page for more information about signals.

### `-X`

Leave the target process stopped on exit from the `exec` system call in order to allow a debugger to attach to it. If you attach `dbx` to the process, use the `dbx` commands `ignore PROF` and `ignore EMT` to ensure that collection signals are passed on to the `collect` command.

### `-y signal[, r]`

Control recording of data with the signal named *signal*. Whenever the signal is delivered to the process, it switches between the paused state, in which no data is recorded, and the recording state, in which data is recorded. Sample points are always recorded, regardless of the state of the switch.

The signal can be specified by the full signal name, by the signal name without the initial letters `SIG`, or by the signal number. Do not use a signal that is used by the program or that would terminate execution. Suggested signals are `SIGUSR1` and `SIGUSR2`. Signals can be delivered to a process by the `kill(1)` command.

If you use both the `-l` and the `-y` options, you must use different signals for each option.

When the `-y` option is used, the Collector is started in the recording state if the optional `r` argument is given, otherwise it is started in the paused state. If the `-y` option is not used, the Collector is started in the recording state.

If you use this option and your program has its own signal handler, you should make sure that the signal that you specify with `-y` is passed on to the Collector's signal handler, and is not intercepted or ignored.

See the `signal(3HEAD)` man page for more information about signals.

## Output Options

`-d` *directory-name*

Place the experiment in directory *directory-name*. This option only applies to individual experiments and not to experiment groups. If the directory does not exist, the `collect` command prints an error message and exits.

`-g` *group-name*

Make the experiment part of experiment group *group-name*. If *group-name* does not end in `.erg`, the `collect` command prints an error message and exits. If the group exists, the experiment is added to it. If *group-name* is not an absolute path, the experiment group is placed in the directory *directory-name* if a directory has been specified with `-d`, otherwise it is placed in the current directory.

`-o` *experiment-name*

Use *experiment-name* as the name of the experiment to be recorded. If *experiment-name* does not end in `.er`, the `collect` command prints an error message and exits. See [“Experiment Names” on page 58](#) for more information on experiment names and how the Collector handles them.

### -A *option*

Control whether or not load objects used by the target process should be archived or copied into the recorded experiment. The allowed values of option are:

- `off` – do not archive load objects into the experiment.
- `on` – archive load objects into the experiment.
- `copy` – copy and archive load objects into the experiment.

If you expect to copy experiments to a different machine from which they were recorded, or to read the experiments from a different machine, you should specify `-A copy`. Using this option does not copy any source files or object files into the experiment. You should ensure that those files are accessible on the machine to which you are copying the experiment.

### -L *size*

Limit the amount of profiling data recorded to *size* megabytes. The limit applies to the sum of the amounts of clock-based profiling data, hardware-counter overflow profiling data, and synchronization wait tracing data, but not to sample points. The limit is only approximate, and can be exceeded.

When the limit is reached, no more profiling data is recorded but the experiment remains open until the target process terminates. If periodic sampling is enabled, sample points continue to be written.

The default limit on the amount of data recorded is 2000 Mbytes. This limit was chosen because the Performance Analyzer cannot process experiments that contain more than 2 Gbytes of data. To remove the limit, set *size* to `unlimited` or `none`.

## Other Options

### -n

Do not run the target but print the details of the experiment that would be generated if the target were run. This is a dry run option.

### -R

Display the text version of the performance tools readme in the terminal window. If the readme is not found, a warning is printed. No further arguments are examined, and no further processing is done.

-V

Print the current version of the `collect` command. No further arguments are examined, and no further processing is done.

-V

Print the current version of the `collect` command and detailed information about the experiment being run.

---

## Collecting Data Using the `dbx` collector Subcommands

To run the Collector from `dbx`:

1. Load your program into `dbx` by typing the following command.

```
% dbx program
```

2. Use the `collector` command to enable data collection, select the data types, and set any optional parameters.

```
(dbx) collector subcommand
```

To get a listing of available `collector` subcommands, type:

```
(dbx) help collector
```

You must use one `collector` command for each subcommand.

3. Set up any `dbx` options you wish to use and run the program.

If a subcommand is incorrectly given, a warning message is printed and the subcommand is ignored. A complete listing of the `collector` subcommands follows.

# Data Collection Subcommands

The following subcommands control the types of data that are collected by the Collector. They are ignored with a warning if an experiment is active.

## `profile option`

Controls the collection of clock-based profiling data. The allowed values for *option* are:

- `on` – Enables clock-based profiling with the default profiling interval of 10 ms.
- `off` – Disables clock-based profiling.
- `timer interval` – Sets the profiling interval. The allowed values of *interval* are
  - `on` – Use the default profiling interval of 10 milliseconds.
  - `lo[w]` – Use the low-resolution profiling interval of 100 milliseconds.
  - `hi[gh]` – Use the high-resolution profiling interval of 1 millisecond. In the Solaris 7 Operating System and earlier versions of the Solaris 8 Operating System, high-resolution profiling must be explicitly enabled. See [“Limitations on Clock-Based Profiling” on page 54](#) for information on enabling high-resolution profiling.
  - `value` – Set the profiling interval to *value*. The default units for *value* are milliseconds. You can specify *value* as an integer or a floating-point number. The numeric value can optionally be followed by the suffix `m` to select millisecond units or `u` to select microsecond units. The value should be a multiple of the clock resolution. If the value is larger than the clock resolution but not a multiple it is rounded down. If the value is smaller than the clock resolution it is set to the clock resolution. In both cases a warning message is printed.

The default setting is 10 ms.

The Collector collects clock-based profiling data by default, unless the collection of hardware-counter overflow profiling data is turned on using the `hwprofile` subcommand.

## `hwprofile option`

Controls the collection of hardware-counter overflow profiling data. If you attempt to enable hardware-counter overflow profiling on systems that do not support it, `dbx` returns a warning message and the command is ignored. The allowed values for *option* are:

- `on` – Turns on hardware-counter overflow profiling. The default action is to collect data for the `cycles` counter at the normal overflow value.

- `off` – Turns off hardware-counter overflow profiling.
- `list` – Returns a list of available counters See [“Hardware Counter Lists” on page 31](#) for a description of the list. If your system does not support hardware-counter overflow profiling, `dbx` returns a warning message.
- `counter name value [ name2 value2 ]` – Selects the hardware counter *name*, and sets its overflow value to *value*; optionally selects a second hardware counter *name2* and sets its overflow value to *value2*. The overflow value can be one of the following.
  - `hi[gh]` – The high-resolution value for the chosen counter is used. The abbreviation `h` is also supported.
  - `lo[w]` – The low-resolution value for the chosen counter is used.
  - *number* – The overflow value. Must be a positive integer.
  - `on` – The default overflow value is used.

The two counters must use different registers. If they do not, a warning message is printed and the command is ignored.

If the hardware counter counts events that relate to memory access, the counter name can be prefixed with a `+` sign to turn on searching for the true PC of the instruction that caused the counter overflow. If the search is successful, the PC and the effective address that was referenced are stored in the event data packet.

The Collector does not collect hardware-counter overflow profiling data by default. If hardware-counter overflow profiling is enabled and a `profile` command has not been given, clock-based profiling is turned off.

See also [“Limitations on Hardware-Counter Overflow Profiling” on page 55](#).

## `synctrace` *option*

Controls the collection of synchronization wait tracing data. The allowed values for *option* are

- `on` – Enable synchronization wait tracing with the default threshold.
- `off` – Disable synchronization wait tracing.
- `threshold value` – Sets the threshold for the minimum synchronization delay. The allowed values for *value* are:
  - `all` – Use a zero threshold. This option forces all synchronization events to be recorded.
  - `calibrate` – Set the threshold value by calibration at runtime. (Equivalent to `on`.)
  - `off` – Turn off synchronization wait tracing.
  - `on` – Use the default threshold, which is to set the value by calibration at runtime. (Equivalent to `calibrate`.)



- *number* – Set the threshold to *number*, given as a positive integer in microseconds. If *value* is 0, all events are traced.

By default, the Collector does not collect synchronization wait tracing data.

### heaptrace *option*

Controls the collection of heap tracing data. The allowed values for *option* are

- *on* – Enables heap tracing.
- *off* – Disables heap tracing.

By default, the Collector does not collect heap tracing data.

### mpitrace *option*

Controls the collection of MPI tracing data. The allowed values for *option* are

- *on* – Enables tracing of MPI calls.
- *off* – Disables tracing of MPI calls.

By default, the Collector does not collect MPI tracing data.

### sample *option*

Controls the sampling mode. The allowed values for *option* are:

- *periodic* – Enables periodic sampling.
- *manual* – Disables periodic sampling. Manual sampling remains enabled.
- *period value* – Sets the sampling interval to *value*, given in seconds.

By default, periodic sampling is enabled, with a sampling interval *value* of 1 second.

### dbxsample { *on* | *off* }

Controls the recording of samples when dbx stops the target process. The meanings of the keywords are as follows:

- *on* – A sample is recorded each time dbx stops the target process.
- *off* – Samples are not recorded when dbx stops the target process.

By default, samples are recorded when dbx stops the target process.

# Experiment Control Subcommands

## `disable`

Disables data collection. If a process is running and collecting data, it terminates the experiment and disables data collection. If a process is running and data collection is disabled, it is ignored with a warning. If no process is running, it disables data collection for subsequent runs.

## `enable`

Enables data collection. If a process is running but data collection is disabled, it enables data collection and starts a new experiment. If a process is running and data collection is disabled, it is ignored with a warning. If no process is running, it enables data collection for subsequent runs.

You can enable and disable data collection as many times as you like during the execution of any process. Each time you enable data collection, a new experiment is created.

## `pause`

Suspends the collection of data, but leaves the experiment open. Sample points are still recorded. This subcommand is ignored if data collection is already paused.

## `resume`

Resumes data collection after a `pause` has been issued. This subcommand is ignored if data is being collected.

## `sample record name`

Record a sample packet with the label *name*. The label is displayed in the Event tab of the Performance Analyzer.

# Output Subcommands

The following subcommands define storage options for the experiment. They are ignored with a warning if an experiment is active.

## archive *mode*

Set the mode for archiving the experiment. The allowed values for *mode* are

- *on* – normal archiving of load objects
- *off* – no archiving of load objects
- *copy* – copy load objects into experiment in addition to normal archiving

If you intend to move the experiment to a different machine, or read it from another machine, you should enable the copying of load objects. If an experiment is active, the command is ignored with a warning. This command does not copy source files or object files into the experiment.

## limit *value*

Limit the amount of profiling data recorded to *value* megabytes. The limit applies to the sum of the amounts of clock-based profiling data, hardware-counter overflow profiling data, and synchronization wait tracing data, but not to sample points. The limit is only approximate, and can be exceeded.

When the limit is reached, no more profiling data is recorded but the experiment remains open and sample points continue to be recorded.

The default limit on the amount of data recorded is 2000 Mbytes. This limit was chosen because the Performance Analyzer cannot process experiments that contain more than 2 Gbytes of data. To remove the limit, set *value* to *unlimited* or *none*.

## store *option*

Governs where the experiment is stored. This command is ignored with a warning if an experiment is active. The allowed values for *option* are:

- *directory directory-name* – Sets the directory where the experiment and any experiment group is stored. This subcommand is ignored with a warning if the directory does not exist.
- *experiment experiment-name* – Sets the name of the experiment. If the experiment name does not end in *.er*, the subcommand is ignored with a warning. See [“Where the Data Is Stored” on page 58](#) for more information on experiment names and how the Collector handles them.
- *group group-name* – Sets the name of the experiment group. If the group name does not end in *.erg*, the subcommand is ignored with a warning. If the group already exists, the experiment is added to the group. If the directory name has been set using the *store directory* subcommand and the group name is not an absolute path, the group name is prefixed with the directory name.

## Information Subcommands

`show`

Shows the current setting of every Collector control.

`status`

- Reports on the status of any open experiment.

---

## Collecting Data From a Running Process

The Collector allows you to collect data from a running process. If the process is already under the control of `dbx` (either in the command line version or in the IDE), you can pause the process and enable data collection using the methods described in previous sections.

---

**Note** – For information on starting the Performance Analyzer from the IDE, see the Performance Analyzer Readme, which is available through the documentation index at `file:/opt/SUNWspr/docs/index.html`. If the Sun Studio 8 software is not installed in the `/opt` directory, ask your system administrator for the equivalent path on your system.

---

If the process is not under the control of `dbx`, you can attach `dbx` to it, collect performance data, and then detach from the process, leaving it to continue. If you want to collect performance data for selected descendant processes, you must attach `dbx` to each process.

To collect data from a running process that is not under the control of `dbx`:

### 1. Determine the program's process ID (PID).

If you started the program from the command line and put it in the background, its PID will be printed to standard output by the shell. Otherwise you can determine the program's PID by typing the following.

```
% ps -ef | grep program-name
```

### 2. Attach to the process.

- From the Debug menu of the IDE, choose Debug → Attach to Solaris Process and select the process using the dialog box. Use the online help for instructions.
- From `dbx`, type the following.

```
(dbx) attach program-name pid
```

If `dbx` is not already running, type the following.

```
% dbx program-name pid
```

See the manual, *Debugging a Program With dbx*, for more details on attaching to a process. Attaching to a running process pauses the process.

### 3. Start data collection.

- From the Debug menu of the IDE, choose Performance Toolkit → Enable Collector and use the dialog box to set up the data collection parameters. Then choose Debug → Continue to resume execution of the process.
- From `dbx`, use the `collector` command to set up the data collection parameters and the `cont` command to resume execution of the process.

### 4. Detach from the process.

When you have finished collecting data, pause the program and then detach the process from `dbx`.

- In the IDE, right-click the session for the process in the Sessions view of the Debugger window and choose Detach from the contextual menu. If the Sessions view is not displayed, click the Sessions button at the top of the Debugger window.
- From `dbx`, type the following.

```
(dbx) detach
```

If you want to collect any kind of tracing data, you must preload the Collector library, `libcollector.so`, before you run your program, because the library provides wrappers to the real functions that enable data collection to take place. In addition, the Collector adds wrapper functions to other system library calls to guarantee the integrity of performance data. If you do not preload the Collector library, these wrapper functions cannot be inserted. See [“Using System Libraries” on page 46](#) for more information on how the Collector interposes on system library functions.

To preload `libcollector.so`, you must set both the name of the library and the path to the library using environment variables. Use the environment variable `LD_PRELOAD` to set the name of the library. Use the environment variables `LD_LIBRARY_PATH`, `LD_LIBRARY_PATH_32`, and/or `LD_LIBRARY_PATH_64` to set the path to the library. (`LD_LIBRARY_PATH` is used if the `_32` and `_64` variants are not defined.) If you have already defined these environment variables, add new values to them.

**TABLE 3-2** Environment Variable Settings for Preloading the Library `libcollector.so`

Environment variable	Value
<code>LD_PRELOAD</code>	<code>libcollector.so</code>
<code>LD_LIBRARY_PATH</code>	<code>/opt/SUNWspro/lib</code>
<code>LD_LIBRARY_PATH_32</code>	<code>/opt/SUNWspro/lib</code>
<code>LD_LIBRARY_PATH_64</code>	<code>/opt/SUNWspro/lib/v9</code>

If your Sun compilers and tools are not installed in `/opt/SUNWspro`, ask your system administrator for the correct path. You can set the full path in `LD_PRELOAD`, but doing this can create complications when using SPARC V9 64-bit architecture.

**Note** – Remove the `LD_PRELOAD` and `LD_LIBRARY_PATH` settings after the run, so they do not remain in effect for other programs that are started from the same shell.

If you want to collect data from an MPI program that is already running, you must attach a separate instance of `dbx` to each process and enable the Collector for each process. When you attach `dbx` to the processes in an MPI job, each process will be halted and restarted at a different time. The time difference could change the interaction between the MPI processes and affect the performance data you collect. To minimize this problem, one solution is to use `pstop(1)` to halt all the processes. However, once you attach `dbx` to the processes, you must restart them from `dbx`, and there will be a timing delay in restarting the processes, which can affect the synchronization of the MPI processes. See also [“Collecting Data From MPI Programs” on page 78](#).

---

## Collecting Data From MPI Programs

The Collector can collect performance data from multi-process programs that use the Sun Message Passing Interface (MPI) library. The MPI library is included in the Sun HPC ClusterTools™ software. You should use the latest version of the ClusterTools

software if possible, which is 4.0, but you can use 3.1 or a compatible version. To start the parallel jobs, use the Sun Cluster Runtime Environment (CRE) command `mprun`. See the Sun HPC ClusterTools documentation for more information. For information about MPI and the MPI standard, see the MPI web site <http://www.mcs.anl.gov/mpi>.

Because of the way MPI and the Collector are implemented, each MPI process records a separate experiment. Each experiment must have a unique name. Where and how the experiment is stored depends on the kinds of file systems that are available to your MPI job. Issues about storing experiments are discussed in the next subsection.

To collect data from MPI jobs, you can either run the `collect` command under MPI or start `dbx` under MPI and use the `dbx collector` subcommands. Each of these options is discussed in subsequent subsections.

## Storing MPI Experiments

Because multiprocessing environments can be complex, there are some issues about storing MPI experiments you should be aware of when you collect performance data from MPI programs. These issues concern the efficiency of data collection and storage, and the naming of experiments. See [“Where the Data Is Stored” on page 58](#) for information on naming experiments, including MPI experiments.

Each MPI process that collects performance data creates its own experiment. When an MPI process creates an experiment, it locks the experiment directory. All other MPI processes must wait until the lock is released before they can use the directory. Thus, if you store the experiments on a file system that is accessible to all MPI processes, the experiments are created sequentially, but if you store the experiments on file systems that are local to each MPI process, the experiments are created concurrently.

If you store the experiments on a common file system and specify an experiment name in the standard format, *experiment.n.ex*, the experiments have unique names. The value of *n* is determined by the order in which MPI processes obtain a lock on the experiment directory, and cannot be guaranteed to correspond to the MPI rank of the process. If you attach `dbx` to MPI processes in a running MPI job, *n* will be determined by the order of attachment.

If you store the experiments on a local file system and specify an experiment name in the standard format, the names are not unique. For example, suppose you ran an MPI job on a machine with 4 single-processor nodes labelled `node0`, `node1`, `node2`

and node3. Each node has a local disk called `/scratch`, and you store the experiments in directory `username` on this disk. The experiments created by the MPI job have the following full path names.

```
node0:/scratch/username/test.1.er
node1:/scratch/username/test.1.er
node2:/scratch/username/test.1.er
node3:/scratch/username/test.1.er
```

The full name including the node name is unique, but in each experiment directory there is an experiment named `test.1.er`. If you move the experiments to a common location after the MPI job is completed, you must make sure that the names remain unique. For example, to move these experiments to your home directory, which is assumed to be accessible from all nodes, and rename the experiments, type the following commands.

```
rsh node0 'er_mv /scratch/username/test.1.er test.0.er'
rsh node1 'er_mv /scratch/username/test.1.er test.1.er'
rsh node2 'er_mv /scratch/username/test.1.er test.2.er'
rsh node3 'er_mv /scratch/username/test.1.er test.3.er'
```

For large MPI jobs, you might want to move the experiments to a common location using a script. Do not use the Unix commands `cp` or `mv`; see [“Manipulating Experiments” on page 149](#) for information on how to copy and move experiments.

If you do not specify an experiment name, the Collector uses the MPI rank to construct an experiment name with the standard form `experiment.n.er`, but in this case `n` is the MPI rank. The stem, `experiment`, is the stem of the experiment group name if you specify an experiment group, otherwise it is `test`. The experiment names are unique, regardless of whether you use a common file system or a local file system. Thus, if you use a local file system to record the experiments and copy them to a common file system, you will not have to rename the experiments when you copy them and reconstruct any experiment group file.

If you do not know which local file systems are available to you, use the `df -lk` command or ask your system administrator. You should always make sure that the experiments are stored in a directory that already exists, that is uniquely defined and that is not in use for any other experiment. You should also make sure that the file system has enough space for the experiments. See [“Estimating Storage Requirements” on page 60](#) for information on how to estimate the space needed.



---

**Note** – If you copy or move experiments between computers or nodes you cannot view the annotated source code or source lines in the annotated disassembly code unless you have access to the load objects and source files that were used to run the experiment, or a copy with the same path and timestamp.

---

## Running the `collect` Command Under MPI

To collect data with the `collect` command under the control of MPI, use the following syntax.

```
% mprun -np n collect [collect-arguments] program-name [program-arguments]
```

Here, *n* is the number of processes to be created by MPI. This procedure creates *n* separate instances of `collect`, each of which records an experiment. Read the section “Where the Data Is Stored” on page 58 for information on where and how to store the experiments.

To ensure that the sets of experiments from different MPI runs are stored separately, you can create an experiment group with the `-g` option for each MPI run. The experiment group should be stored on a file system that is accessible to all MPI processes. Creating an experiment group also makes it easier to load the set of experiments for a single MPI run into the Performance Analyzer. An alternative to creating a group is to specify a separate directory for each MPI run with the `-d` option.

## Collecting Data by Starting `dbx` Under MPI

To start `dbx` and collect data under the control of MPI, use the following syntax.

```
% mprun -np n dbx program-name < collection-script
```

Here, *n* is the number of processes to be created by MPI and *collection-script* is a `dbx` script that contains the commands necessary to set up and start data collection. This procedure creates *n* separate instances of `dbx`, each of which records an experiment on one of the MPI processes. If you do not define the experiment name, the experiment will be labelled with the MPI rank. Read the section “Storing MPI Experiments” on page 79 for information on where and how to store the experiments.

You can name the experiments with the MPI rank by using the collection script and a call to `MPI_Comm_rank()` in your program. For example, in a C program you would insert the following line.

```
ier = MPI_Comm_rank(MPI_COMM_WORLD, &me);
```

In a Fortran program you would insert the following line.

```
call MPI_Comm_rank(MPI_COMM_WORLD, me, ier)
```

If this call was inserted at line 17, for example, you could use a script like this.

```
stop at 18
run program-arguments
rank=${me}
collector enable
collector store filename experiment.$rank.er
cont
quit
```

---

## Using collect With ppgsz

You can use `collect` with `ppgsz(1)` by running `collect` on the `ppgsz` command and specifying the `-F` on flag. The founder experiment is on the `ppgsz` executable and uninteresting. If your path finds the 32-bit version of `ppgsz`, and the experiment is run on a system that supports 64-bit processes, the first thing it will do is `exec` its 64-bit version, creating `_x1.er`. That executable forks, creating `_x1_f1.er`.

The child process attempts to `exec` the named target in the first directory on your path, then in the second, and so forth, until one of the `execs` succeed. If, for example, the third attempt succeeds, the first two descendant experiments are named `_x1_f1_x1.er` and `_x1_f1_x2.er`, and both are completely empty. The experiment on the target is the one from the successful `exec`, the third one in the example, and is named `_x1_f1_x3.er`, stored under the founder experiment. It can be processed directly by invoking the Analyzer or `er_print` on `test.1.er/_x1_f1_x3.er`.

If the 64-bit `ppgsz` is the initial process, or if the 32-bit `ppgsz` is invoked on a 32-bit kernel, the fork child that `execs` the real target has its data in `_f1.er`, and the real target's experiment is in `_f1_x3.er`, assuming the same path properties as in the example above.

# The `er_print` Command Line Performance Analysis Tool

---

This chapter explains how to use the `er_print` utility for performance analysis. The `er_print` utility prints an ASCII version of the various displays supported by the Performance Analyzer. The information is written to standard output unless you redirect it to a file or printer. You must give the `er_print` utility the name of one or more experiments or experiment groups generated by the Collector as arguments. You can use the `er_print` utility to display the performance metrics for functions, for callers and callees; the source code and disassembly listings; sampling information; data-space data; and execution statistics.

This chapter covers the following topics.

- [er\\_print Syntax](#)
- [Metric Lists](#)
- [Commands That Control the Function List](#)
- [Commands That Control the Callers-Callees List](#)
- [Commands That Control the Leak and Allocation Lists](#)
- [Commands That Control the Source and Disassembly Listings](#)
- [Commands That Control the Data Space List](#)
- [Commands That List Experiments, Samples, Threads, and LWPs](#)
- [Commands That Control Selections](#)
- [Commands That Control Load Object Selection](#)
- [Commands That List Metrics](#)
- [Commands That Control Output](#)
- [Commands That Print Other Displays](#)
- [Commands That Set Defaults](#)
- [Commands That Set Defaults Only For the Performance Analyzer](#)
- [Miscellaneous Commands](#)
- [Examples](#)

For a description of the data collected by the Collector, see [Chapter 2](#).

For instructions on how to use the Performance Analyzer to display information in a graphical format, see the online help.

---

## er\_print Syntax

The command-line syntax for `er_print` is as follows.

```
er_print [ -script script | -command | - | -v ] experiment-list
```

The options for `er_print` are listed in [TABLE 4-1](#).

**TABLE 4-1** Options for the `er_print` Command

Option	Description
-	Read <code>er_print</code> commands entered from the keyboard.
-script <i>script</i>	Read commands from the file <i>script</i> , which contains a list of <code>er_print</code> commands, one per line. If the <code>-script</code> option is not present, <code>er_print</code> reads commands from the terminal or from the command line.
-command [ <i>argument</i> ]	Process the given command.
-v	Display version information and exit.

Multiple options can appear on the `er_print` command line. They are processed in the order they appear. You can mix scripts, hyphens, and explicit commands in any order. The default action if you do not supply any commands or scripts is to enter interactive mode, in which commands are entered from the keyboard. To exit interactive mode type **quit** or Ctrl-D.

The commands accepted by `er_print` are listed in the following sections. You can abbreviate any command with a shorter string as long as the command is unambiguous.

---

## Metric Lists

Many of the `er_print` commands use a list of metric keywords. The syntax of the list is as follows.

```
metric-keyword-1 [ : metric-keyword2... ]
```

Except for the `size`, `address`, and `name` keywords, a metric keyword consists of three parts: a metric type string, a metric visibility string, and a metric name string. These are joined with no spaces, as follows.

```
<type><visibility><name>
```

The metric type and metric visibility strings are composed of type and visibility characters.

The allowed metric type characters are given in [TABLE 4-2](#). A metric keyword that contains more than one type character is expanded into a list of metric keywords. For example, `ie.user` is expanded into `i.user:e.user`.

**TABLE 4-2** Metric Type Characters

Character	Description
e	Show exclusive metric value
i	Show inclusive metric value
a	Show attributed metric value (only for callers-callees metrics)

The allowed metric visibility characters are given in [TABLE 4-3](#). The order of the visibility characters in the visibility string does not matter: it does not affect the order in which the corresponding metrics are displayed. For example, both `i%.user` and `i.%user` are interpreted as `i.user:i%user`.

Metrics that differ only in the visibility are always displayed together in the standard order. If two metric keywords that differ only in the visibility are separated by some other keywords, the metrics appear in the standard order at the position of the first of the two metrics.

**TABLE 4-3** Metric Visibility Characters

Character	Description
.	Show metric as a time. Applies to timing metrics and hardware counter metrics that measure cycle counts. Interpreted as "+" for other metrics.
%	Show metric as a percentage of the total program metric. For attributed metrics in the callers-callees list, show metric as a percentage of the inclusive metric for the selected function.
+	Show metric as an absolute value. For hardware counters, this value is the event count. Interpreted as "." for timing metrics.
!	Do not show any metric value. Cannot be used in combination with other visibility characters.

When both type and visibility strings have more than one character, the type is expanded first. Thus `i.e.%user` is expanded to `i.e.%user:e.%user`, which is then interpreted as `i.user:i%user:e.user:e%user`.

The visibility characters period (`.`), plus (`+`), and percent sign (`%`), are equivalent for the purposes of defining the sort order. Thus `sort i%user`, `sort i.user`, and `sort i+user` all mean that the Analyzer should sort by inclusive user CPU time if it is visible in any form, and `sort i!user` means the Analyzer should sort by inclusive user CPU time, whether or not it is visible.

[TABLE 4-4](#) lists the available `er_print` metric name strings for timing metrics, synchronization delay metrics, memory allocation metrics, MPI tracing metrics, and the two common hardware counter metrics. For other hardware counter metrics, the metric name string is the same as the counter name. A list of counter names can be obtained by using the `collect` command with no arguments. See [“Hardware-Counter Overflow Profiling Data” on page 30](#) for more information on hardware counters.

**TABLE 4-4** Metric Name Strings

Category	String	Description
<b>Timing metrics</b>	<code>user</code>	User CPU time
	<code>wall</code>	Wall-clock time
	<code>total</code>	Total LWP time
	<code>system</code>	System CPU time
	<code>wait</code>	CPU wait time
	<code>unlock</code>	User lock time
	<code>text</code>	Text-page fault time
	<code>data</code>	Data-page fault time
	<code>owait</code>	Other wait time
<b>Synchronization delay metrics</b>	<code>sync</code>	Synchronization wait time
	<code>syncn</code>	Synchronization wait count
<b>MPI tracing metrics</b>	<code>mpitime</code>	Time spent in MPI calls
	<code>mpisend</code>	Number of MPI send operations
	<code>mpibytessent</code>	Number of bytes sent in MPI send operations
	<code>mpireceive</code>	Number of MPI receive operations
	<code>mpibytesrecv</code>	Number of bytes received in MPI receive operations
	<code>mpiother</code>	Number of calls to other MPI functions

**TABLE 4-4** Metric Name Strings (*Continued*)

Category	String	Description
<b>Memory allocation metrics</b>	<code>alloc</code>	Number of allocations
	<code>balloc</code>	Bytes allocated
	<code>leak</code>	Number of leaks
	<code>bleak</code>	Bytes leaked
<b>Hardware counter overflow metrics</b>	<code>cycles</code>	CPU cycles
	<code>insts</code>	Instructions issued

In addition to the name strings listed in [TABLE 4-4](#), there are two name strings that can only be used in default metrics lists. These are `hwc`, which matches any hardware counter name, and `any`, which matches any metric name string. Also note that `cycles` and `insts` are common to SPARC® and Intel, but other flavors also exist that are architecture-specific. To list all available counters, use the `collect` command with no arguments.

---

## Commands That Control the Function List

The following commands control how the function information is displayed.

`functions`

Write the function list with the currently selected metrics. The function list includes all functions in load objects that are selected for display of functions, and any load objects whose functions are hidden with the `object_select` command.

The number of lines written can be limited by using the `limit` command (see [“Commands That Control Output” on page 100](#)).

The default metrics printed are exclusive and inclusive user CPU time, in both seconds and percentage of total program metric. You can change the current metrics displayed with the `metrics` command. This must be done before you issue the `functions` command. You can also change the defaults with the `dmetrics` command.

For applications written in the Java programming language, the displayed function information varies depending on whether Java mode is set to `on`, `expert`, or `off`. When Java mode is set to `on`, the displayed function information includes metrics against the Java methods, and any native methods called. Setting the Java mode to `expert` shows HotSpot-compiled methods separately from the interpreted version of the method. Setting the mode to `off` shows functions from the JVM itself, rather than from the Java application being interpreted by the JVM, along with any compiled methods and native methods.

## `metrics` *metric\_spec*

Specify a selection of function-list metrics. The string *metric\_spec* can either be the keyword `default`, which restores the default metric selection, or a list of metric keywords, separated by colons. The following example illustrates a metric list.

```
% metrics i.user:i%user:e.user:e%user
```

This command instructs `er_print` to display the following metrics:

- Inclusive user CPU time in seconds
- Inclusive user CPU time percentage
- Exclusive user CPU time in seconds
- Exclusive user CPU time percentage

When the `metrics` command is finished, a message is printed showing the current metric selection. For the preceding example the message is as follows.

```
current: i.user:i%user:e.user:e%user:name
```

For information on the syntax of metric lists, see [“Metric Lists” on page 84](#). To see a listing of the available metrics, use the `metric_list` command.

If a `metrics` command has an error in it, it is ignored with a warning, and the previous settings remain in effect.

## `sort` *metric\_spec*

Sort the function list on the specified metric. The string *metric\_spec* is one of the metric keywords described in [“Metric Lists” on page 84](#), as shown in this example.

```
% sort i.user
```



This command tells `er_print` to sort the function list by inclusive user CPU time. If the metric is not in the experiments that have been loaded, a warning is printed and the command is ignored. When the command is finished, the sort metric is printed.

## `fsummary`

Write a summary metrics panel for each function in the function list. The number of panels written can be limited by using the `limit` command (see [“Commands That Control Output” on page 100](#)).

The summary metrics panel includes the name, address and size of the function or load object, and for functions, the name of the source file, object file and load object, and all the recorded metrics for the selected function or load object, both exclusive and inclusive, as values and percentages.

## `fsingle function_name [N]`

Write a summary metrics panel for the specified function. The optional parameter `N` is needed for those cases where there are several functions with the same name. The summary metrics panel is written for the `N`th function with the given function name. When the command is given on the command line, `N` is required; if it is not needed it is ignored. When the command is given interactively without `N` but `N` is required, a list of functions with the corresponding `N` value is printed.

For a description of the summary metrics for a function, see the `fsummary` command description.

---

# Commands That Control the Callers-Callees List

The following commands control how the caller and callee information is displayed.

## callers-callees

Print the callers-callees panel for each of the functions, in the order in which they are sorted. The number of panels written can be limited by using the `limit` command (see [“Commands That Control Output” on page 100](#)). The selected (center) function is marked with an asterisk, as shown in this example.

Attr.	Excl.	Incl.	Name
User CPU	User CPU	User CPU	
sec.	sec.	sec.	
4.440	0.	42.910	commandline
0.	0.	4.440	*gpf
4.080	0.	4.080	gpf_b
0.360	0.	0.360	gpf_a

In this example, `gpf` is the selected function; it is called by `commandline`, and it calls `gpf_a` and `gpf_b`.

## `cmetrics` *metric\_spec*

Specify a selection of callers-callees metrics. *metric\_spec* is a list of metric keywords, separated by colons, as shown in this example.

```
% cmetrics i:user:i%user:a:user:a%user
```

This command instructs `er_print` to display the following metrics.

- Inclusive user CPU time in seconds
- Inclusive user CPU time percentage
- Attributed user CPU time in seconds
- Attributed user CPU time percentage

When the `cmetrics` command is finished, a message is printed showing the current metric selection. For the preceding example the message is as follows.

```
current: i:user:i%user:a:user:a%user:name
```

For information on the syntax of metric lists, see [“Metric Lists” on page 84](#). To see a listing of the available metrics, use the `cmetric_list` command.

```
csingle function_name [N]
```

Write the callers-callees panel for the named function. The optional parameter *N* is needed for those cases where there are several functions with the same name. The callers-callees panel is written for the *N*th function with the given function name. When the command is given on the command line, *N* is required; if it is not needed it is ignored. When the command is given interactively without *N* but *N* is required, a list of functions with the corresponding *N* value is printed.

```
csort metric_spec
```

Sort the callers-callees display by the specified metric. The string *metric\_spec* is one of the metric keywords described in [“Metric Lists” on page 84](#), as shown in this example.

```
% csort a.user
```

This command tells `er_print` to sort the callers-callees display by attributed user CPU time. When the command finishes, the sort metric is printed.

---

## Commands That Control the Leak and Allocation Lists

This section describes the commands that relate to memory allocations and deallocations.

```
leaks
```

Display a list of memory leaks, aggregated by common call stack. Each entry presents the total number of leaks and the total bytes leaked for the given call stack. The list is sorted by the number of bytes leaked.

```
allocs
```

Display a list of memory allocations, aggregated by common call stack. Each entry presents the number of allocations and the total bytes allocated for the given call stack. The list is sorted by the number of bytes allocated.

---

# Commands That Control the Source and Disassembly Listings

The following commands control how annotated source and disassembly code is displayed.

## `pcs`

Write a list of program counters (PCs) and their metrics, ordered by the current sort metric. The list includes lines that show aggregated metrics for each load object whose functions are hidden with the `object_select` command.

## `psummary`

Write the summary metrics panel for each PC in the PC list, in the order specified by the current sort metric.

## `lines`

Write a list of source lines and their metrics, ordered by the current sort metric. The list includes lines that show aggregated metrics for each function that does not have line-number information, or whose source file is unknown, and lines that show aggregated metrics for each load object whose functions are hidden with the `object_select` command.

## `lsummary`

Write the summary metrics panel for each line in the lines list, in the order specified by the current sort metric.

`source { filename | function_name } [N]`

Write out annotated source code for either the specified file or the file containing the specified function. The file in either case must be in a directory in your path.

Use the optional parameter *N* (a positive integer) only in those cases where the file or function name is ambiguous; in this case, the *N*th possible choice is used. If you give an ambiguous name without the numeric specifier, `er_print` prints a list of possible object-file names; if the name you gave was a function, the name of the function is appended to the object-file name, and the number that represents the value of *N* for that object file is also printed.

```
disasm { filename | function_name } [N]
```

Write out annotated disassembly code for either the specified file, or the file containing the specified function. The file in either case must be in a directory in your path.

The optional parameter *N* is used in the same way as for the `source` command.

```
scc com_spec
```

Specify the classes of compiler commentary that are shown in the annotated source listing. The class list is a colon-separated list of classes, containing zero or more of the following message classes.

**TABLE 4-5** Compiler Commentary Message Classes

Class	Meaning
b[asic]	Show the basic level messages.
v[ersion]	Show version messages, including source file name and last modified date, versions of the compiler components, compilation date and options.
pa[rallel]	Show messages about parallelization.
q[uey]	Show questions about the code that affect its optimization.
l[oop]	Show messages about loop optimizations and transformations.
pi[pe]	Show messages about pipelining of loops.
i[nline]	Show messages about inlining of functions.
m[emops]	Show messages about memory operations, such as load, store, prefetch.
f[e]	Show front-end messages.
all	Show all messages.
none	Do not show any messages.

The classes `all` and `none` cannot be used with other classes.

If no `scc` command is given, the default class shown is `basic`. If the `scc` command is given with an empty *class-list*, compiler commentary is turned off. The `scc` command is normally used only in a `.er.rc` file.

### `sthresh value`

Specify the threshold percentage for highlighting metrics in the annotated source code. If the value of any metric is equal to or greater than *value* % of the maximum value of that metric for any source line in the file, the line on which the metrics occur have `##` inserted at the beginning of the line.

### `dcc com_spec`

Specify the classes of compiler commentary that are shown in the annotated disassembly listing. The class list is a colon-separated list of classes. The list of available classes is the same as the list of classes for annotated source code listing. The following options can be added to the class list.

**TABLE 4-6** Additional Options for the `dcc` Command

Option	Meaning
<code>h[ex]</code>	Show the hexadecimal value of the instructions.
<code>noh[ex]</code>	Do not show the hexadecimal value of the instructions.
<code>s[rc]</code>	Interleave the source listing in the annotated disassembly listing.
<code>nos[rc]</code>	Do not interleave the source listing in the annotated disassembly listing.
<code>as[rc]</code>	Interleave the annotated source code in the annotated disassembly listing.

### `dthresh value`

Specify the threshold percentage for highlighting metrics in the annotated disassembly code. If the value of any metric is equal to or greater than *value* % of the maximum value of that metric for any instruction line in the file, the line on which the metrics occur have `##` inserted at the beginning of the line.

`setpath path_list`

Set the path used to find source, object, etc. files. *path\_list* is a colon-separated list of directories. If any directory has a colon character in it, it should be escaped with a backslash. The special directory name, *\$expts*, refers to the set of current experiments, in the order in which they were loaded; it may be abbreviated with a single *\$* character.

The default setting is: *\$expts:..*. The compiled-in full pathname will be used if a file is not found in searching the current path setting.

`setpath` with no argument prints the current path.

`addpath path_list`

Append *path\_list* to the current `setpath` settings.

---

## Commands That Control the Data Space List

`data_objects`

Write the list of data objects with their metrics. Applicable only to HW counter experiments where aggressive backtracking was specified, and for objects in files that were compiled with `-xhwcprof`. (Available on SPARC for C only). See the *C User's Guide* or the `cc(1)` man page for further information.

`data_osingle name [N]`

Write the summary metrics panel for the named data object. The optional parameter *N* is needed for those cases where the object name is ambiguous. When the directive is on the command-line, *N* is required; if it is not needed, it is ignored. Applicable only to HW counter experiments where aggressive backtracking was specified, and for objects in files that were compiled with `-xhwcprof`. (Available on SPARC for C only). See the *C User's Guide* or the `cc(1)` man page for further information.

## data\_olayout

Write the annotated data object layouts for all program data objects with data-derived metric data, in the order they are defined in the experiment's load objects. Each aggregate data object is shown with the total metrics attributed to it, followed by all of its elements in offset order, each with their own metrics and an indicator of its size and location relative to 32-byte blocks.

---

# Commands That List Experiments, Samples, Threads, and LWPs

This section describes the commands that list experiments, samples, threads, and LWPs.

## experiment\_list

Display the full list of experiments loaded with their ID number. Each experiment is listed with an index, which is used when selecting samples, threads, or LWPs.

The following example is an example of an experiment list.

```
(er_print) experiment_list
ID Experiment
== =====
1 test.1.er
2 test.6.er
```

## sample\_list

Display the list of samples currently selected for analysis.

The following example is an example of a sample list.

```
(er_print) sample_list
Exp Sel      Total
=== ===== =====
1 1-6         31
2 7-10,15    31
```



`lwp_list`

Display the list of LWPs currently selected for analysis.

`thread_list`

Display the list of threads currently selected for analysis.

`cpu_list`

Display the list of CPUs currently selected for analysis.

---

## Commands That Control Selections

### Selection Lists

The syntax of a selection is shown in the following example. This syntax is used in the command descriptions.

```
[experiment-list : ] selection-list [+ [experiment-list : ] selection-list ... ]
```

Each selection list can be preceded by an experiment list, separated from it by a colon and no spaces. You can make multiple selections by joining selection lists with a + sign.

The experiment list and the selection list have the same syntax, which is either the keyword `all` or a list of numbers or ranges of numbers ( $n-m$ ) separated by commas but no spaces, as shown in this example.

```
2, 4, 9-11, 23-32, 38, 40
```

The experiment numbers can be determined by using the `exp_list` command.

Some examples of selections are as follows.

```
1:1-4+2:5, 6  
all:1, 3-6
```

In the first example, objects 1 through 4 are selected from experiment 1 and objects 5 and 6 are selected from experiment 2. In the second example, objects 1 and 3 through 6 are selected from all experiments. The objects may be LWPs, threads, or samples.

## Selection Commands

The commands to select LWPs, samples, CPUs, and threads are not independent. If the experiment list for a command is different from that for the previous command, the experiment list from the latest command is applied to all three selection targets – LWPs, samples, and threads, in the following way.

- Existing selections for experiments not in the latest experiment list are turned off.
- Existing selections for experiments in the latest experiment list are kept.
- Selections are set to `all` for targets for which no selection has been made.

`sample_select` *sample\_spec*

Select the samples for which you want to display information. The list of samples you selected is displayed when the command finishes.

`lwp_select` *lwp\_spec*

Select the LWPs about which you want to display information. The list of LWPs you selected is displayed when the command finishes.

`thread_select` *thread\_spec*

Select the threads about which you want to display information. The list of threads you selected is displayed when the command finishes.

`cpu_select` *cpu\_spec*

Select the CPUs about which you want to display information. The list of CPUs you selected is displayed when the command finishes.

---

# Commands That Control Load Object Selection

`object_list`

Display the list of load objects. The name of each load object is preceded either by a `yes` which indicates that the functions of that object are shown in the function list, or by a `no` which indicates that the functions of that object are not shown in the function list.

The following is an example of a load object list.

```
(er_print) object_list
Sel Load Object
=== =====
yes /tmp/var/synprog/synprog
yes /opt/SUNWspro/lib/libcollector.so
yes /usr/lib/libdl.so.1
yes /usr/lib/libc.so.1
```

`object_select` *object1,object2,...*

Select the load objects for which you want to display information about the functions in the load object. *object-list* is a list of load objects, separated by commas but no spaces. For load objects that are not selected, information for the entire load object is displayed instead of information for the functions in the load object.

The names of the load objects should be either full path names or the basename. If an object name itself contains a comma, you must surround the name with double quotation marks.

---

# Commands That List Metrics

The following commands list the currently selected metrics and all available metric keywords.

`metric_list`

Display the currently selected metrics in the function list and a list of metric keywords that you can use in other commands (for example, `metrics` and `sort`) to reference various types of metrics in the function list.

`cmetric_list`

Display the currently selected metrics in the callers-callees list and a list of metric keywords that you can use in other commands (for example, `cmetrics` and `csort`) to reference various types of metrics in the callers-callees list.

---

**Note** – Attributed metrics can only be specified for display with the `cmetrics` command, not the `metrics` command, and displayed only with the `callers-callees` command, not the `functions` command.

---

## Commands That Control Output

The following commands control `er_print` display output.

`outfile { filename | - }`

Close any open output file, then open *filename* for subsequent output. If you specify a dash (-) instead of *filename*, output is written to standard output.

`limit n`

Limit output to the first *n* entries of the report; *n* is an unsigned positive integer.

`name { long | short }`

Specify whether to use the long or the short form of function names (C++ only).

```
javamode { on | expert | off }
```

Set the mode for Java experiments to on (show the Java model), expert (show the Java model, but show HotSpot-compiled methods independently from interpreted methods), or off (show the machine model).

---

## Commands That Print Other Displays

`header exp_id`

Display descriptive information about the specified experiment. The *exp\_id* can be obtained from the `exp_list` command. If the *exp\_id* is `all` or is not given, the information is displayed for all experiments loaded.

Following each header, any errors or warnings are printed. Headers for each experiment are separated by a line of dashes.

*exp\_id* is required on the command line, but not in a script or in interactive mode.

`objects`

List the load objects with any error or warning messages that result from the use of the load object for performance analysis. The number of load objects listed can be limited by using the `limit` command (see [“Commands That Control Output” on page 100](#)).

`overview exp_id`

Write out the sample data of each of the currently selected samples for the specified experiment. The *exp\_id* can be obtained from the `exp_list` command. If the *exp\_id* is `all` or is not given, the sample data is displayed for all experiments. *exp\_id* is required on the command line, but not in a script or in interactive mode.

`statistics exp_id`

Write out execution statistics, aggregated over the current sample set for the specified experiment. For information on the definitions and meanings of the execution statistics that are presented, see the `getrusage(3C)` and `proc(4)` man pages. The execution statistics include statistics from system threads for which the

Collector does not collect any data. The standard threads library in the Solaris 7 and Solaris 8 Operating System creates system threads that are not profiled. These threads spend most of their time sleeping, and the time shows in the statistics display as Other Wait time.

The *exp\_id* can be obtained from the `experiment_list` command. If the *exp\_id* is not given, the sum of data for all experiments is displayed, aggregated over the sample set for each experiment. If *exp\_id* is `all`, the sum and the individual statistics for each experiment are displayed.

---

## Commands That Set Defaults

You can use the following commands to set the defaults for `er_print` and for the Performance Analyzer. You can only use these commands for setting defaults: they cannot be used in input for `er_print`. They can be included in a defaults file named `.er.rc`. Some of the commands only apply to the Performance Analyzer.

A defaults file can be included in your home directory, to set defaults for all experiments, or in any other directory, to set defaults locally. When `er_print`, `er_src` or the Performance Analyzer is started, the current directory and your home directory are scanned for defaults files, which are read if they are present, and the system defaults file is also read. Defaults from the `.er.rc` file in your home directory override the system defaults, and defaults from the `.er.rc` file in the current directory override both home and system defaults.

---

**Note** – To ensure that you read the defaults file from the directory where your experiment is stored, you must start the Performance Analyzer or the `er_print` utility from that directory.

---

The defaults file can also include the `scc`, `sthresh`, `dcc`, and `dthresh` commands. Multiple `dmetrics` and `dsort` commands can be given in a defaults file, and the commands within a file are concatenated.

### `dmetrics` *metric\_spec*

Specify the default metrics to be displayed or printed in the function list. The syntax and use of the metric list is described in the section [“Metric Lists” on page 84](#). The order of the metric keywords in the list determines the order in which the metrics are presented and the order in which they appear in the Metric chooser in the Performance Analyzer.

Default metrics for the Callers-Callees list are derived from the function list default metrics by adding the corresponding attributed metric before the first occurrence of each metric name in the list.

`dsort metric_spec`

Specify the default metric by which the function list is sorted. The sort metric is the first metric in this list that matches a metric in any loaded experiment, subject to the following conditions:

- If the entry in *metric\_spec* has a visibility string of an exclamation point, !, the first metric whose name matches is used, whether it is visible or not.
- If the entry in *metric\_spec* has any other visibility string, the first visible metric whose name matches is used.

The syntax and use of the metric list is described in the section [“Metric Lists” on page 84](#).

The default sort metric for the Callers-Callees list is the attributed metric corresponding to the default sort metric for the function list.

`gdemangle library.so`

Set the path to the shared object that supports an API to demangle C++ function names. The shared object must export the C function `cplus_demangle()`, conforming to the GNU standard `libiberty.so` interface.

---

# Commands That Set Defaults Only For the Performance Analyzer

`tlmode` *tl\_mode*

Set the display mode options for the Timeline tab of the Performance Analyzer. The list of options is a colon-separated list. The allowed options are described in the following table.

**TABLE 4-7** Timeline Display Mode Options

Option	Meaning
<code>lw[p]</code>	Display events for LWPs
<code>t[hread]</code>	Display events for threads
<code>c[pu]</code>	Display events for CPUs
<code>r[oot]</code>	Align call stack at the root
<code>le[af]</code>	Align call stack at the leaf
<code>d[epth]</code> <i>nn</i>	Set the maximum depth of the call stack that can be displayed

The options `lwp`, `thread`, and `cpu` are mutually exclusive, as are `root` and `leaf`. If more than one of a set of mutually exclusive options is included in the list, the last one is the only one that is used.

`tldata` *tl\_data*

Select the default data types shown in the Timeline tab of the Performance Analyzer. The types in the type list are separated by colons. The allowed types are listed in the following table.

**TABLE 4-8** Timeline Display Data Types

Type	Meaning
<code>sa[mple]</code>	Display sample data
<code>c[lock]</code>	Display clock profiling data
<code>hw[c]</code>	Display hardware counter profiling data



**TABLE 4-8** Timeline Display Data Types (*Continued*)

Type	Meaning
sy[nctrace]	Display thread synchronization tracing data
mp[itrace]	Display MPI tracing data
he[aptrace]	Display heap tracing data

`datamode { on | off }`

Set the mode for showing dataspace-related screens to on (tabs are visible), or off (do not have them visible).

---

## Miscellaneous Commands

`mapfile load-object { mapfilename | - }`

Write a mapfile for the specified load object to the file *mapfilename*. If you specify a dash (-) instead of *mapfilename*, `er_print` writes the mapfile to standard output.

`script file`

Process additional commands from the script file *file*.

`version`

Print the current release number of `er_print`.

`quit`

Terminate processing of the current script, or exit interactive mode.

`help`

Print a list of `er_print` commands.

---

# Examples

- The following example generates a gprof-like list from an experiment. The output is a file called `er_print.out` which lists the top 100 functions, followed by caller-callee data, sorted by attributed user time for each.

```
er_print -outfile er_print.out -metrics e.user:e%user\  
-sort e.user -limit 100 -functions -cmetrics a.user:a%user\  
-csort a.user -callers-callees test.1.er
```

You can also simplify this example into the following independent commands. However, keep in mind that each call to `er_print` in a large experiment or application can be time intensive:

- ```
er_print -metrics e.user:e%user -sort e.user \  
-limit 100 -functions test.1.er
```
- ```
er_print -cmetrics a.user:a%user -csort a.user \  
-callers-callees test.1.er
```
- This example summarizes how time is spent in functions.  

```
er_print -functions test.*.er
```
- This example shows the caller-callee relationships.  

```
er_print -callers-callees test.*.er
```
- This example shows which source lines are hot. Source-line information assumes the code was compiled and linked with `-g`. Append a trailing underscore to the function name for Fortran functions and routines. The 1 after the function name is used to distinguish between multiple instances of the *myfunction*.  

```
er_print -source myfunction 1 test.*.er
```
- This example shows only the compiler commentary. It is not necessary to run your program in order to use this command.  

```
er_src -myfile.o
```
- These examples uses wall-clock profiling to list functions and callers-callees.  

```
er_print -metrics ei.%wall -functions test.*.er  
er_print -cmetrics aei.%wall -callers-callees test.*.er
```
- This example shows high-level MPI functions. There are lots of internal software layers in MPI, but this is one way to see just the entry points. There may be some duplicates of symbols but you can ignore the duplicates.  

```
er_print -functions test.*.er | grep PMPI_
```

# Understanding the Performance Analyzer and Its Data

---

The Performance Analyzer reads the event data that is collected by the Collector and converts it into performance metrics. The metrics are computed for various elements in the structure of the target program, such as instructions, source lines, functions, and load objects. In addition to a header, the data recorded for each event collected has two parts:

- Some event-specific data that is used to compute metrics
- A call stack of the application that is used to associate those metrics with the program structure

The process of associating the metrics with the program structure is not always straightforward, due to the insertions, transformations, and optimizations made by the compiler. This chapter describes the process and discusses the effect on what you see in the Performance Analyzer displays.

This chapter covers the following topics:

- How Data Collection Works
- [Interpreting Performance Metrics](#)
- [Call Stacks and Program Execution](#)
- [Mapping Addresses to Program Structure](#)
- [Mapping Data Addresses to Program Data Objects](#)
- [Annotated Code Listings](#)

---

## How Data Collection Works

The output from a data collection run is an experiment, which is stored as a directory with various internal files and subdirectories in the file system.

# Experiment Format

All experiments must have three files:

- A log file; an ASCII file that contains information about what data was collected, what versions of various components, and a record of various events during the life of the target.
- A loadobjects file; an ASCII file that records the time-dependent information about what loadobjects are loaded into the address space of the target, and the times at which they are loaded or unloaded.
- An overview file; a binary file containing usage information recorded at every sample point in the experiment.

In addition, experiments have binary data files representing the profile events in the life of the process. Each data file has a series of events, as described below under [“Interpreting Performance Metrics” on page 111](#). Separate files are used for each type of data, but each file is shared by all LWPs in the target. The data files are named as follows:

**TABLE 5-1** Data Types and Corresponding File Names

Data Type	File Name
clock-profiling	profile
HWC-profiling	hwcounters
Synchronization-tracing	synctrace
Heap-tracing	heaptrace
MPI-tracing	mpitrace

For clock-profiling, or HW-counter-overflow-profiling, the data is written in a signal handler invoked by the clock-tick or counter-overflow. For synchronization tracing, heap tracing, or MPI tracing, data is written from `libcollector.so` routines that are interposed by `LD_PRELOAD` on the normal user-invoked routines. Each such interposition routine partially fills in a data record, then invokes the normal user-invoked routine, and fills in the rest of the data record when that routine returns, and writes the record to the data file.

All data files are memory-mapped, and filled in blocks. The records are filled in such a way as to always have a valid record structure, so that experiments can be read as they are being written. The buffer management strategy is designed to minimize contention and serialization between LWPs.

## The archive subdirectory

Each experiment has an archive subdirectory, which contains binary files describing each load object referenced in the loadobjects file. These files are produced by `er_archive`, which runs at the end of data collection. If the process terminates abnormally, `er_archive` may not be invoked, in which case, the archive files are written by `er_print` or Analyzer when first invoked on the experiment.

## Descendant Processes

Descendant processes write their experiments into subdirectories within the founder-process' experiment. These subdirectories are named with an underscore, a code letter (f for fork and x for exec), and a number are added to its immediate creator's experiment name, giving the genealogy of the descendant. For example, if the experiment name for the founder process is `test.1.er`, the experiment for the child process created by its third fork is `test.1.er/_f3.er`. If that child process execs a new image, the corresponding experiment name is `test.1.er/_f3_x1.er`. Descendant experiments consist of the same files as the parent experiment, but they do not have descendant experiments (all descendants are represented by subdirectories in the founder experiment), and they do not have archive subdirectories (all archiving is done into the founder experiment).

## Dynamic Functions

Experiments where the target creates dynamic functions have additional records in the loadobjects file describing those functions, and an additional file, `dyntext`, containing a copy of the actual instructions of the dynamic functions. The copy is needed to produce annotated disassembly of dynamic functions.

## Java Experiments

Java experiments also have additional records in the loadobjects file, both for dynamic functions created by the JVM for its internal purposes, and for dynamically-compiled (HotSpot) versions of the target Java methods.

Java experiments also have additional records in the loadobjects file, both for dynamic functions created by the JVM for its internal purposes, and for dynamically-compiled (HotSpot) versions of the target Java methods.

In addition, Java experiments have a `JAVA_CLASSES` file, containing information about all of the user's Java classes invoked.

Java heap- and synchronization tracing data is recorded using a JVMPI agent, which is part of `libcollector.so`. It receives events that are mapped into the recorded trace events. The agent also receives events for class loading and HotSpot compilation, that are used to write the `JAVA_CLASSES` file, and the Java-compiled method records in the `loadobjects` file.

## Recording Experiments

There are three different ways to record an experiment, with the `collect` command, with `dbx` creating a process, and with `dbx` creating an experiment from a running process.

### `collect` Experiments

When `collect` is used to record an experiment, the `collect` program itself creates the experiment directory and sets `LD_PRELOAD` to ensure that `libcollector.so` is preloaded into the target's address space. It then sets environment variables to inform `libcollector` about the experiment name, and data collection options, and `exec's` the target on top of itself.

`libcollector.so` is responsible for writing all experiment files.

### `dbx` Experiments, Creating the Process

When `dbx` is used to launch a process with data collection enabled, it also creates the experiment directory, and ensures preloading of `libcollector.so`. It stops the process at a breakpoint before its first instruction, and then calls an initialization routine in `libcollector` to start the data collection.

Java experiments can not be collected by `dbx`, since `dbx` uses a JVMDI agent for debugging, and that agent can not coexist with the JVMPI agent needed for data collection.

### `dbx` Experiments, on a Running Process

When `dbx` is used to start an experiment on a running process, it creates the experiment directory, but can not use `LD_PRELOAD`. It makes an interactive function call into the target to `dlopen libcollector.so`, and then calls `libcollector.so's` initialization routine, just as it does when creating the process. Data is written by `libcollector.so` just as in a `collect` experiment.

Since `libcollector.so` was not in the target address space when the process started, any data collection that depends on interposition on user-callable functions (synchronization tracing, heap tracing, MPI tracing), may not work. In general, the symbols will have already been resolved to the underlying functions, so the interposition can not happen. Furthermore, the following of descendant processes also depends on interposition, and will not work properly for experiments created by `dbx` on a running process.

If the user has explicitly `LD_PRELOAD`'ed `libcollector.so` before starting the process with `dbx`, or before using `dbx` to attach to the running process, tracing data may be collected.

---

## Interpreting Performance Metrics

The data for each event contains a high-resolution timestamp, a thread ID, an LWP ID, and a processor ID. The first three of these can be used to filter the metrics in the Performance Analyzer by time, thread or LWP. See the `getcpuid(2)` man page for information on processor IDs. On systems where `getcpuid` is not available, the processor ID is -1, which maps to Unknown.

In addition to the common data, each event generates specific raw data, which is described in the following sections. Each section also contains a discussion of the accuracy of the metrics derived from the raw data and the effect of data collection on the metrics.

### Clock-Based Profiling

The event-specific data for clock-based profiling consists of an array of profiling interval counts for each of the ten microstates maintained by the kernel for each LWP. At the end of the profiling interval, the count for the microstate of each LWP is incremented by 1, and a profiling signal is scheduled. The array is only recorded and reset when the LWP is in user mode in the CPU. If the LWP is in user mode when the profiling signal is scheduled, the array element for the User-CPU state is 1, and the array elements for all the other states are 0. If the LWP is not in user mode, the data is recorded when the LWP next enters user mode, and the array can contain an accumulation of counts for various states.

The call stack is recorded at the same time as the data. If the LWP is not in user mode at the end of the profiling interval, the call stack cannot change until the LWP enters user mode again. Thus the call stack always accurately records the position of the program counter at the end of each profiling interval.

The metrics to which each of the microstates contributes are shown in [TABLE 5-2](#).

**TABLE 5-2** How Kernel Microstates Contribute to Metrics

Kernel Microstate	Description	Metric Name
LMS_USER	Running in user mode	User CPU Time
LMS_SYSTEM	Running in system call or page fault	System CPU Time
LMS_TRAP	Running in any other trap	System CPU Time
LMS_TFAULT	Asleep in user text page fault	Text Page Fault Time
LMS_DFAULT	Asleep in user data page fault	Data Page Fault Time
LMS_KFAULT	Asleep in kernel page fault	Other Wait Time
LMS_USER_LOCK	Asleep waiting for user-mode lock	User Lock Time
LMS_SLEEP	Asleep for any other reason	Other Wait Time
LMS_STOPPED	Stopped ( <code>/proc</code> , job control, or <code>lwp_stop</code> )	Other Wait Time
LMS_WAIT_CPU	Waiting for CPU	Wait CPU Time

## Accuracy of Timing Metrics

Timing data is collected on a statistical basis, and is therefore subject to all the errors of any statistical sampling method. For very short runs, in which only a small number of profile packets is recorded, the call stacks might not represent the parts of the program which consume the most resources. You should run your program for long enough or enough times to accumulate hundreds of profile packets for any function or source line you are interested in.

In addition to statistical sampling errors, there are specific errors that arise from the way the data is collected and attributed and the way the program progresses through the system. Some of the circumstances in which inaccuracies or distortions can appear in the timing metrics are described in what follows.

- When an LWP is created, the time it has spent before the first profile packet is recorded is less than the profiling interval, but the entire profiling interval is ascribed to the microstate recorded in the first profile packet. If there are many LWPs created the error can be many times the profiling interval.
- When an LWP is destroyed, some time is spent after the last profile packet is recorded. If there are many LWPs destroyed the error can be many times the profiling interval.
- LWP rescheduling can occur during a profiling interval. As a consequence, the recorded state of the LWP might not represent the microstate in which it spent most of the profiling interval. The errors are likely to be larger when there are more LWPs to run than there are processors to run them.



- It is possible for a program to behave in a way which is correlated with the system clock. In this case, the profiling interval always expires when the LWP is in a state which might represent a small fraction of the time spent, and the call stacks recorded for a particular part of the program are overrepresented. On a multiprocessor system, it is possible for the profiling signal to induce a correlation: processors that are interrupted by the profiling signal while they are running LWPs for the program are likely to be in the Trap-CPU microstate when the microstate is recorded.
- The kernel records the microstate value when the profiling interval expires. When the system is under heavy load, that value might not represent the true state of the process. This situation is likely to result in overaccounting of the Trap-CPU or Wait-CPU microstate.
- The threads library sometimes discards profiling signals when it is in a critical section, resulting in an underaccounting of timing metrics.
- When the system clock is being synchronized with an external source, the time stamps recorded in profile packets do not reflect the profiling interval but include any adjustment that was made to the clock. The clock adjustment can make it appear that profile packets are lost. The time period involved is usually several seconds, and the adjustments are made in increments.

In addition to the inaccuracies just described, timing metrics are distorted by the process of collecting data. The time spent recording profile packets never appears in the metrics for the program, because the recording is initiated by profiling signal. (This is another instance of correlation.) The user CPU time spent in the recording process is distributed over whatever microstates are recorded. The result is an underaccounting of the User CPU Time metric and an overaccounting of other metrics. The amount of time spent recording data is typically less than one percent of the CPU time for the default profiling interval.

## Comparisons of Timing Metrics

If you compare timing metrics obtained from the profiling done in a clock-based experiment with times obtained by other means, you should be aware of the following issues.

For a single-threaded application, the total LWP time recorded for a process is usually accurate to a few tenths of a percent, compared with the values returned by `gethrtime(3C)` for the same process. The CPU time can vary by several percentage points from the values returned by `gethrtime(3C)` for the same process. Under heavy load, the variation might be even more pronounced. However, the CPU time differences do not represent a systematic distortion, and the relative times reported for different functions, source-lines, and such are not substantially distorted.

For multithreaded applications using unbound threads, differences in values returned by `gethrvtime()` could be meaningless. This is because `gethrvtime()` returns values for an LWP, and a thread can change from one LWP to another.

The LWP times that are reported in the Performance Analyzer can differ substantially from the times that are reported by `vmstat`, because `vmstat` reports times that are summed over CPUs. If the target process has more LWPs than the system on which it is running has CPUs, the Performance Analyzer shows more wait time than `vmstat` reports.

The microstate timings that appear in the Statistics tab of the Performance Analyzer and the `er_print` statistics display are based on process file system usage reports, for which the times spent in the microstates are recorded to high accuracy. See the `proc(4)` man page for more information. You can compare these timings with the metrics for the `<Total>` function, which represents the program as a whole, to gain an indication of the accuracy of the aggregated timing metrics. However, the values displayed in the Statistics tab can include other contributions that are not included in the timing metric values for `<Total>`. These contributions come from the following sources:

- Threads that are created by the system that are not profiled. The standard threads library in the Solaris Operating System versions 7 and 8 creates system threads that are not profiled. These threads spend most of their time sleeping, and the time shows in the Statistics tab as Other Wait time.
- Periods of time in which data collection is paused.

## Synchronization Wait Tracing

The Collector collects synchronization delay events by tracing calls to the functions in the threads library, `libthread.so`, or to the real time extensions library, `librt.so`. The event-specific data consists of high-resolution timestamps for the request and the grant (beginning and end of the call that is traced), and the address of the synchronization object (the mutex lock being requested, for example). The thread and LWP IDs are the IDs at the time the data is recorded. The wait time is the difference between the request time and the grant time. Only events for which the wait time exceeds the specified threshold are recorded. The synchronization wait tracing data is recorded in the experiment at the time of the grant.

If the program uses bound threads, the LWP on which the waiting thread is scheduled cannot perform any other work until the event that caused the delay is completed. The time spent waiting appears both as Synchronization Wait Time and as User Lock Time. User Lock Time can be larger than Synchronization Wait Time because the synchronization delay threshold screens out delays of short duration.

If the program uses unbound threads, it is possible for the LWP on which the waiting thread is scheduled to have other threads scheduled on it and continue to perform user work. The User Lock Time is zero if all LWPs are kept busy while some threads are waiting for a synchronization event. However, the Synchronization Wait Time is not zero because it is associated with a particular thread, not with the LWP on which the thread is running.

The wait time is distorted by the overhead for data collection. The overhead is proportional to the number of events collected. The fraction of the wait time spent in overhead can be minimized by increasing the threshold for recording events.

## Hardware-Counter Overflow Profiling

Hardware-counter overflow profiling data includes a counter ID and the overflow value. The value can be larger than the value at which the counter is set to overflow, because the processor executes some instructions between the overflow and the recording of the event. This is especially true of cycle and instruction counters, which are incremented much more frequently than counters such as floating-point operations or cache misses. The delay in recording the event also means that the program counter address recorded with call stack does not correspond exactly to the overflow event.

See [“Attribution of Hardware Counter Overflows” on page 147](#) for more information. See also the discussion of [“Traps” on page 119](#). Traps and trap handlers can cause significant differences between reported User CPU time and time reported by the cycle counter.

The amount of data collected depends on the overflow value. Choosing a value that is too small can have the following consequences.

- The amount of time spent collecting data can be a substantial fraction of the execution time of the program. The collection run might spend most of its time handling overflows and writing data instead of running the program.
- A substantial fraction of the counts can come from the collection process. These counts are attributed to the collector function `collector_record_counters`. If you see high counts for this function, the overflow value is too small.
- The collection of data can alter the behavior of the program. For example, if you are collecting data on cache misses, the majority of the misses could come from flushing the collector instructions and profiling data from the cache and replacing it with the program instructions and data. The program would appear to have a lot of cache misses, but without data collection there might in fact be very few cache misses.

Choosing a value that is too large can result in too few overflows for good statistics. The counts that are accrued after the last overflow are attributed to the collector function `collector_final_counters`. If you see a substantial fraction of the counts in this function, the overflow value is too large.

## Heap Tracing

The Collector records tracing data for calls to the memory allocation and deallocation functions `malloc`, `realloc`, `memalign` and `free` by interposing on these functions. If your program bypasses these functions to allocate memory, tracing data is not recorded. Tracing data is not recorded for Java memory management, which uses a different mechanism.

The functions that are traced could be loaded from any of a number of libraries. The data that you see in the Performance Analyzer might depend on the library from which a given function is loaded.

If a program makes a large number of calls to the traced functions in a short space of time, the time taken to execute the program can be significantly lengthened. The extra time is used in recording the tracing data.

## MPI Tracing

MPI tracing records information about calls to MPI library functions. The event-specific data consists of high-resolution timestamps for the request and the grant (beginning and end of the call that is traced), the number of send and receive operations and the number of bytes sent or received. Tracing is done by interposing on the calls to the MPI library. The interposing functions do not have detailed information about the optimization of data transmission, nor about transmission errors, so the information that is presented represents a simple model of the data transmission, which is explained in the following paragraphs.

The number of bytes received is the length of the buffer as defined in the call to the MPI function. The actual number of bytes received is not available to the interposing function.

Some of the Global Communication functions have a single origin or a single receiving process known as the root. The accounting for such functions is done as follows:

- Root sends data to all processes, itself included.
- Root receives data from all processes, itself included.
- Each process communicates with each process, itself included

The following examples illustrate the accounting procedure. In these examples, G is the size of the group.

For a call to `MPI_Bcast()`,

- Root sends G packets of N bytes, one packet to each process, including itself
- All G processes in the group (including root) receive N bytes

For a call to `MPI_Allreduce()`,

- Each process sends G packets of N bytes
- Each process receives G packets of N bytes

For a call to `MPI_Reduce_scatter()`,

- Each process sends G packets of N/G bytes
- Each process receives G packets of N/G bytes

---

## Call Stacks and Program Execution

A call stack is a series of program counter addresses (PCs) representing instructions from within the program. The first PC, called the leaf PC, is at the bottom of the stack, and is the address of the next instruction to be executed. The next PC is the address of the call to the function containing the leaf PC; the next PC is the address of the call to that function, and so forth, until the top of the stack is reached. Each such address is known as a return address. The process of recording a call stack involves obtaining the return addresses from the program stack and is referred to as unwinding the stack. For information on unwind failures, see [“Incomplete Stack Unwinds” on page 128](#).

The leaf PC in a call stack is used to assign exclusive metrics from the performance data to the function in which that PC is located. Each PC on the stack, including the leaf PC, is used to assign inclusive metrics to the function in which it is located.

Most of the time, the PCs in the recorded call stack correspond in a natural way to functions as they appear in the source code of the program, and the Performance Analyzer’s reported metrics correspond directly to those functions. Sometimes, however, the actual execution of the program does not correspond to a simple intuitive model of how the program would execute, and the Performance Analyzer’s reported metrics might be confusing. See [“Mapping Addresses to Program Structure” on page 130](#) for more information about such cases.

# Single-Threaded Execution and Function Calls

The simplest case of program execution is that of a single-threaded program calling functions within its own load object.

When a program is loaded into memory to begin execution, a context is established for it that includes the initial address to be executed, an initial register set, and a stack (a region of memory used for scratch data and for keeping track of how functions call each other). The initial address is always at the beginning of the function `_start()`, which is built into every executable.

When the program runs, instructions are executed in sequence until a branch instruction is encountered, which among other things could represent a function call or a conditional statement. At the branch point, control is transferred to the address given by the target of the branch, and execution proceeds from there. (Usually the next instruction after the branch is already committed for execution: this instruction is called the branch delay slot instruction. However, some branch instructions annul the execution of the branch delay slot instruction.)

When the instruction sequence that represents a call is executed, the return address is put into a register, and execution proceeds at the first instruction of the function being called.

In most cases, somewhere in the first few instructions of the called function, a new frame (a region of memory used to store information about the function) is pushed onto the stack, and the return address is put into that frame. The register used for the return address can then be used when the called function itself calls another function. When the function is about to return, it pops its frame from the stack, and control returns to the address from which the function was called.

## Function Calls Between Shared Objects

When a function in one shared object calls a function in another shared object, the execution is more complicated than in a simple call to a function within the program. Each shared object contains a Program Linkage Table, or PLT, which contains entries for every function external to that shared object that is referenced from it. Initially the address for each external function in the PLT is actually an address within `ld.so`, the dynamic linker. The first time such a function is called, control is transferred to the dynamic linker, which resolves the call to the real external function and patches the PLT address for subsequent calls.

If a profiling event occurs during the execution of one of the three PLT instructions, the PLT PCs are deleted, and exclusive time is attributed to the call instruction. If a profiling event occurs during the first call through a PLT entry, but the leaf PC is not one of the PLT instructions, any PCs that arise from the PLT and code in `ld.so` are replaced by a call to an artificial function, `@plt`, which accumulates inclusive time.

There is one such artificial function for each shared object. If the program uses the `LD_AUDIT` interface, the PLT entries might never be patched, and non-leaf PCs from `@plt` can occur more frequently.

## Signals

When a signal is sent to a process, various register and stack operations occur that make it look as though the leaf PC at the time of the signal is the return address for a call to a system function, `sigacthandler()`. `sigacthandler()` calls the user-specified signal handler just as any function would call another.

The Performance Analyzer treats the frames resulting from signal delivery as ordinary frames. The user code at the point at which the signal was delivered is shown as calling the system function `sigacthandler()`, and it in turn is shown as calling the user's signal handler. Inclusive metrics from both `sigacthandler()` and any user signal handler, and any other functions they call, appear as inclusive metrics for the interrupted function.

The Collector interposes on `sigaction()` to ensure that its handlers are the primary handlers for the `SIGPROF` signal when clock data is collected and `SIGEMT` signal when hardware counter data is collected.

## Traps

Traps can be issued by an instruction or by the hardware, and are caught by a trap handler. System traps are traps which are initiated from an instruction and trap into the kernel. All system calls are implemented using trap instructions, for example. Some examples of hardware traps are those issued from the floating point unit when it is unable to complete an instruction (such as the `fitos` instruction on the UltraSPARC® III platform), or when the instruction is not implemented in the hardware.

When a trap is issued, the LWP enters system mode. The microstate is usually switched from User CPU state to Trap state then to System state. The time spent handling the trap can show as a combination of System CPU time and User CPU time, depending on the point at which the microstate is switched. The time is attributed to the instruction in the user's code from which the trap was initiated (or to the system call).

For some system calls, it is considered critical to provide as efficient handling of the call as possible. The traps generated by these calls are known as *fast traps*. Among the system functions which generate fast traps are `gethrtime` and `gethrvtime`. In these functions, the microstate is not switched because of the overhead involved.

In other circumstances it is also considered critical to provide as efficient handling of the trap as possible. Some examples of these are TLB (translation lookaside buffer) misses and register window spills and fills, for which the microstate is not switched.

In both cases, the time spent is recorded as User CPU time. However, the hardware counters are turned off because the mode has been switched to system mode. The time spent handling these traps can therefore be estimated by taking the difference between User CPU time and Cycles time, preferably recorded in the same experiment.

There is one case in which the trap handler switches back to user mode, and that is the misaligned memory reference trap for an 8-byte integer which is aligned on a 4-byte boundary in Fortran. A frame for the trap handler appears on the stack, and a call to the handler can appear in the Performance Analyzer, attributed to the integer load or store instruction.

When an instruction traps into the kernel, the instruction following the trapping instruction appears to take a long time, because it cannot start until the kernel has finished executing the trapping instruction.

## Tail-Call Optimization

The compiler can do one particular optimization whenever the last thing a particular function does is to call another function. Rather than generating a new frame, the callee re-uses the frame from the caller, and the return address for the callee is copied from the caller. The motivation for this optimization is to reduce the size of the stack, and, on SPARC® platforms, to reduce the use of register windows.

Suppose that the call sequence in your program source looks like this:

```
A -> B -> C -> D
```

When B and C are tail-call optimized, the call stack looks as if function A calls functions B, C, and D directly.

```
A -> B
```

```
A -> C
```

```
A -> D
```

That is, the call tree is flattened. When code is compiled with the `-g` option, tail-call optimization takes place only at a compiler optimization level of 4 or higher. When code is compiled without the `-g` option, tail-call optimization takes place at a compiler optimization level of 2 or higher.



# Explicit Multithreading

A simple program executes in a single thread, on a single LWP (light-weight process). Multithreaded executables make calls to a thread creation function, to which the target function for execution is passed. When the target exits, the thread is destroyed by the threads library. Newly-created threads begin execution at a function called `_thread_start()`, which calls the function passed in the thread creation call. For any call stack involving the target as executed by this thread, the top of the stack is `_thread_start()`, and there is no connection to the caller of the thread creation function. Inclusive metrics associated with the created thread therefore only propagate up as far as `_thread_start()` and the `<Total>` function.

In addition to creating the threads, the threads library also creates LWPs to execute the threads. Threading can be done either with bound threads, where each thread is bound to a specific LWP, or with unbound threads, where each thread can be scheduled on a different LWP at different times.

- If bound threads are used, the threads library creates one LWP per thread.
- If unbound threads are used, the threads library decides how many LWPs to create to run efficiently, and which LWPs to schedule the threads on. The threads library can create more LWPs at a later time if they are needed. Unbound threads are not part of the Solaris 9 Operating System or of the alternate threads library in the Solaris 8 Operating System.

As an example of the scheduling of unbound threads, when a thread is at a synchronization barrier such as a `mutex_lock`, the threads library can schedule a different thread on the LWP on which the first thread was executing. The time spent waiting for the lock by the thread that is at the barrier appears in the Synchronization Wait Time metric, but since the LWP is not idle, the time is not accrued into the User Lock Time metric.

In addition to the user threads, the standard threads library in the Solaris Operating System versions 7 and 8 creates some threads are used to perform signal handling and other tasks. If the program uses bound threads, additional LWPs are also created for these threads. Performance data is not collected or displayed for these threads, which spend most of their time sleeping. However, the time spent in these threads is included in the process statistics and in the times recorded in the sample data. The threads library in the Solaris 9 Operating System and the alternate threads library in the Solaris 8 Operating System do not create these extra threads.

# Overview of Java Technology-Based Software Execution

To the typical developer, a Java™ technology-based application runs just like any other program. It begins at a main entry point, typically named `class.main`, which may call other methods, just as a C or C++ application does.

To the operating system, an application written in the Java programming language, (pure or mixed with C/C++), runs as a process instantiating the <sup>1</sup>Java virtual machine (JVM). The JVM™ software is compiled from C++ sources and starts execution at `_start`, which calls `main`, and so forth. It reads bytecode from `.class` and/or `.jar` files, and performs the operations specified in that program. Among the operations that can be specified is the dynamic loading of a shared object, and calls into various functions or methods contained within that object.

During execution of a Java technology-based application, most methods are interpreted by the JVM software; these methods are referred to in this document as *interpreted methods*. Other methods may be dynamically compiled by the Java HotSpot™ virtual machine, and are referred to as *compiled methods*. Dynamically compiled methods are loaded into the data space of the application, and may be unloaded at some later point in time. For any particular method, there will be an interpreted version, and there may also be one or more compiled versions. Code written in the Java programming language may also call directly into native-compiled code, either C, C++, or native-compiled (SBA SPARC Bytecode Accelerator) Java; the targets of such calls are referred to as *native methods*.

The JVM software does a number of things that are typically not done by applications written in traditional languages. At startup, it creates a number of regions of dynamically-generated code in its data space. One of these is the actual interpreter code used to process the application's bytecode methods.

During the interpretive execution, the Java HotSpot virtual machine monitors performance, and may decide to take one or more methods that it has been interpreting, generate machine code for them, and execute the more-efficient machine code version, rather than interpret the original. That generated machine code is also in the data space of the process. In addition, other code is generated in the data space to execute the transitions between interpreted and compiled code.

Applications written in the Java programming language are inherently multithreaded, and have one JVM software thread for each thread in the user's program. It also has several housekeeping threads used for signal handling, memory management, and Java HotSpot virtual machine compilation. Depending on the version of `libthread.so` used, there may be a one-to-one correspondence between

---

1. The terms "Java virtual machine" and "JVM" mean a virtual machine for the Java™ platform.

threads and LWPs, or a more complex relationship. But for any version of the library, at any instant a thread may be unscheduled, or scheduled onto an LWP. Data for a thread is not collected while that thread is not scheduled onto an LWP.

The performance tools collect their data by recording events in the life of the each LWP of the process, along with the callstack at the time of the event. At any point in the execution of any application, Java technology-based or otherwise, the callstack represents where the program is in its execution, and how it got there. One important way that mixed-model Java technology-based applications differ from traditional C, C++, and Fortran applications is that at any instant during the run of the target there are two callstacks that are meaningful: a Java callstack, and a machine callstack. Both stacks are collected and correlated with each other.

## Java Processing Representations

There are three representations for displaying performance data for applications written in the Java programming language: the Java representation, the Expert-Java representation, and the Machine representation. The Java Representation is shown by default where the data supports it. The following section summarizes the main differences between these three representations.

### The Java Representation

The Java representation shows compiled and interpreted Java methods by name, and shows native methods in their natural form. During execution, there may be many instances of a particular Java method executed: the interpreted version, and, perhaps, one or more compiled versions. In the Java representation all methods are shown aggregated as a single method. This mode is selected in the analyzer by default.

### The Expert-Java Representation

The Expert-Java representation is similar to the Java Representation, except that some details of the JVM internals that are suppressed in the Java Representation are exposed in the Expert-Java Representation.

### The Machine Representation

The Machine representation shows functions from the JVM itself, rather than from the application being interpreted by the JVM. It also shows all compiled and native methods. The machine representation looks the same as that of applications written

in traditional languages. The callstack shows JVM frames, native frames, and compiled-method frames. Some of the JVM frames represent transition code between interpreted Java, compiled Java, and native code.

## Parallel Execution and Compiler-Generated Body Functions

If your code contains Sun, Cray, or OpenMP parallelization directives, it can be compiled for parallel execution. OpenMP is a feature available with the Sun compilers and tools. Refer to the *OpenMP API User's Guide* and the relevant sections in the *Fortran Programming Guide* and *C User's Guide*, or visit the web site defining the OpenMP standard, <http://www.openmp.org>.

When a loop or other parallel construct is compiled for parallel execution, the compiler-generated code is executed by multiple threads, coordinated by the microtasking library. Parallelization by the Sun compilers follows the procedure outlined below.

### Generation of Body Functions

When the compiler encounters a parallel construct, it sets up the code for parallel execution by placing the body of the construct in a separate *body function* and replacing the construct with a call to a microtasking library function. The microtasking library function is responsible for dispatching threads to execute the body function. The address of the body function is passed to the microtasking library function as an argument.

If the parallel construct is delimited with one of the directives in the following list, then the construct is replaced with a call to the microtasking library function `__mt_MasterFunction_()`.

- The Sun Fortran directive `!$par doall`
- The Cray Fortran directive `c$mic doall`
- A Fortran OpenMP `!$omp PARALLEL`, `!$omp PARALLEL DO`, or `!$omp PARALLEL SECTIONS` directive
- A C or C++ OpenMP `#pragma omp parallel`, `#pragma omp parallel for`, or `#pragma omp parallel sections` directive

A loop that is parallelized automatically by the compiler is also replaced by a call to `__mt_MasterFunction_()`.

If an OpenMP parallel construct contains one or more worksharing do, for or sections directives, each worksharing construct is replaced by a call to the microtasking library function `__mt_Worksharing_()` and a new body function is created for each.

The compiler assigns names to body functions that encode the type of parallel construct, the name of the function from which the construct was extracted, the line number of the beginning of the construct in the original source, and the sequence number of the parallel construct. These mangled names vary from release to release of the microtasking library.

## Parallel Execution Sequence

The program begins execution with only one thread, the main thread. The first time the program calls `__mt_MasterFunction_()`, this function calls the Solaris threads library function, `thr_create()` to create worker threads. Each worker thread executes the microtasking library function `__mt_SlaveFunction_()`, which was passed as an argument to `thr_create()`.

In addition to worker threads, the standard threads library in the Solaris Operating System versions 7 and 8 creates some threads to perform signal handling and other tasks. Performance data is not collected for these threads, which spend most of their time sleeping. However, the time spent in these threads is included in the process statistics and the times recorded in the sample data. The threads library in the Solaris 9 Operating System and the alternate threads library in the Solaris 8 Operating System do not create these extra threads.

Once the threads have been created, `__mt_MasterFunction_()` manages the distribution of available work among the main thread and the worker threads. If work is not available, `__mt_SlaveFunction_()` calls `__mt_WaitForWork_()`, in which the worker thread waits for available work. As soon as work becomes available, the thread returns to `__mt_SlaveFunction_()`.

When work is available, each thread executes a call to `__mt_run_my_job_()`, to which information about the body function is passed. The sequence of execution from this point depends on whether the body function was generated from a parallel sections directive, a parallel do (or parallel for) directive, a parallel workshare directive, or a parallel directive.

- In the parallel sections case, `__mt_run_my_job_()` calls the body function directly.
- In the parallel do or for case, `__mt_run_my_job_()` calls other functions, which depend on the nature of the loop, and the other functions call the body function.
- In the parallel case, `__mt_run_my_job_()` calls the body function directly, and all threads execute the code in the body function until they encounter a call to `__mt_WorkSharing_()`. In this function there is another call to `__mt_run_my_job_()`, which calls the worksharing body function, either

directly in the case of a worksharing section, or through other library functions in the case of a worksharing do or for. If `nowait` was specified in the worksharing directive, each thread returns to the parallel body function and continues executing. Otherwise, the threads return to `__mt_WorkSharing_()`, which calls `__mt_EndOfTaskBarrier_()` to synchronize the threads before continuing.



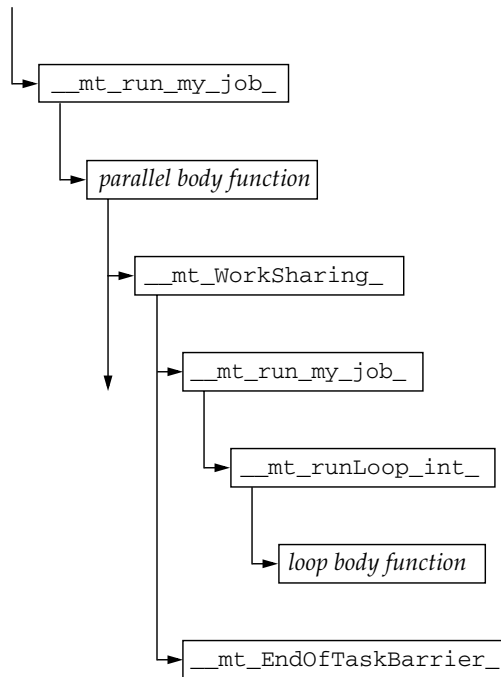
**FIGURE 5-1** Schematic Call Tree for a Multithreaded Program That Contains a Parallel Do or Parallel For Construct

When all parallel work is finished, the threads return to either `__mt_MasterFunction_()` or `__mt_SlaveFunction_()` and call `__mt_EndOfTaskBarrier_()` to perform any synchronization work involved in the termination of the parallel construct. The worker threads then call `__mt_WaitForWork_()` again, while the main thread continues to execute in the serial region.

The call sequence described here applies not only to a program running in parallel, but also to a program compiled for parallelization but running on a single-CPU machine, or on a multiprocessor machine using only one LWP.

The call sequence for a simple parallel do construct is illustrated in [FIGURE 5-1](#). The call stack for a worker thread begins with the threads library function `_thread_start()`, the function which actually calls `__mt_SlaveFunction_()`. The dotted arrow indicates the initiation of the thread as a consequence of a call from `__mt_MasterFunction_()` to `thr_create()`. The continuing arrows indicate that there might be other function calls which are not represented here.

The call sequence for a parallel region in which there is a worksharing do construct is illustrated in [FIGURE 5-2](#). The caller of `__mt_run_my_job_()` is either `__mt_MasterFunction_()` or `__mt_SlaveFunction_()`. The entire diagram can replace the call to `__mt_run_my_job_()` in [FIGURE 5-1](#).



**FIGURE 5-2** Schematic Call Tree for a Parallel Region With a Worksharing Do or Worksharing For Construct

In these call sequences, all the compiler-generated body functions are called from the same function (or functions) in the microtasking library, which makes it difficult to associate the metrics from the body function with the original user function. The Performance Analyzer inserts an imputed call to the body function from the original user function, and the microtasking library inserts an imputed call from the body

function to the barrier function, `__mt_EndOfTaskBarrier()`. The metrics due to the synchronization are therefore attributed to the body function, and the metrics for the body function are attributed to the original function. With these insertions, inclusive metrics from the body function propagate directly to the original function rather than through the microtasking library functions. The side effect of these imputed calls is that the body function appears as a callee of both the original user function and the microtasking functions. In addition, the user function appears to have microtasking library functions as its callers, and can appear to call itself. Double-counting of inclusive metrics is avoided by the mechanism used for recursive function calls (see [“How Recursion Affects Function-Level Metrics” on page 41](#)).

Worker threads typically use CPU time while they are in `__mt_WaitForWork()` in order to reduce latency when new work arrives, that is, when the main thread reaches a new parallel construct. This is known as a busy-wait. However, you can set an environment variable to specify a sleep wait, which shows up in the Performance Analyzer as Other Wait time instead of User CPU time. There are generally two situations where the worker threads spend time waiting for work, where you might want to redesign your program to reduce the waiting:

- When the main thread is executing in a serial region and there is nothing for the worker threads to do
- When the work load is unbalanced, and some threads have finished and are waiting while others are still executing

By default, the microtasking library uses threads that are bound to LWPs. You can override this default in the Solaris Operating System versions 7 and 8 by setting the environment variable `MT_BIND_LWP` to `FALSE`.

---

**Note** – The multiprocessing dispatch process is implementation-dependent and might change from release to release.

---

## Incomplete Stack Unwinds

Stack unwind may fail for a number of reasons:

- If the stack has been corrupted by the user code; if so, the program may core dump, or the data collection code may core dump, depending on exactly how the stack was corrupted.
- If the user code does not follow the standard ABI conventions for function calls. In particular, on the SPARC platform, if the return register, `%o7`, is altered before a save instruction is executed.

On any platform, hand-written assembler code may violate the conventions.



- On the Intel platform, if C or Fortran code is compiled at high optimization, which means that it will not have frame pointers to assist in the unwind.
- On the Intel platform, if C++ code is compiled at any optimization level with the `-noex` or `-features=no%except` options.
- If the leaf PC is in a function after the callee's frame is popped from the stack, but before the function returns.
- If the call stack contains more than about 250 frames, the Collector does not have the space to completely unwind the call stack. In this case, PCs for functions from `_start` to some point in the call stack are not recorded in the experiment, and `<Total>` appears as the caller of the last function whose PC was recorded.

## Intermediate Files

If you generate intermediate files using the `-E` or `-P` compiler options, the Performance Analyzer uses the intermediate file for annotated source code, not the original source file. The `#line` directives generated with `-E` can cause problems in the assignment of metrics to source lines.

The following line appears in annotated source if there are instructions from a function that do not have line numbers referring to the source file that was compiled to generate the function:

```
function_name -- <instructions without line numbers>
```

Line numbers can be absent under the following circumstances:

- You compiled without specifying the `-g` option.
- The debugging information was stripped after compilation, or the executables or object files that contain the information are moved or deleted or subsequently modified.
- The function contains code that was generated from `#include` files rather than from the original source file.
- At high optimization, if code was inlined from a function in a different file.
- The source file has `#line` directives referring to some other file; compiling with the `-E` option, and then compiling the resulting `.i` file is one way in which this happens. It may also happen when you compile with the `-P` flag.
- The object file can not be found to read line number information.
- The file was compiled without the `-g` flag, or the file was stripped.
- The compiler used generates incomplete line number tables.

---

# Mapping Addresses to Program Structure

Once a call stack is processed into PC values, the Performance Analyzer maps those PCs to shared objects, functions, source lines, and disassembly lines (instructions) in the program. This section describes those mappings.

## The Process Image

When a program is run, a process is instantiated from the executable for that program. The process has a number of regions in its address space, some of which are text and represent executable instructions, and some of which are data which is not normally executed. PCs as recorded in the call stack normally correspond to addresses within one of the text segments of the program.

The first text section in a process derives from the executable itself. Others correspond to shared objects that are loaded with the executable, either at the time the process is started, or dynamically loaded by the process. The PCs in a call stack are resolved based on the executable and shared objects loaded at the time the call stack was recorded. Executables and shared objects are very similar, and are collectively referred to as load objects.

Because shared objects can be loaded and unloaded in the course of program execution, any given PC might correspond to different functions at different times during the run. In addition, different PCs might correspond to the same function, when a shared object is unloaded and then reloaded at a different address.

## Load Objects and Functions

Each load object, whether an executable or a shared object, contains a text section with the instructions generated by the compiler, a data section for data, and various symbol tables. All load objects must contain an ELF symbol table, which gives the names and addresses of all the globally-known functions in that object. Load objects compiled with the `-g` option contain additional symbolic information, which can augment the ELF symbol table and provide information about functions that are not global, additional information about object modules from which the functions came, and line number information relating addresses to source lines.

The term *function* is used to describe a set of instructions that represent a high-level operation described in the source code. The term covers subroutines as used in Fortran, methods as used in C++ and Java, and the like. Functions are described cleanly in the source code, and normally their names appear in the symbol table representing a set of addresses; if the program counter is within that set, the program is executing within that function.

In principle, any address within the text segment of a load object can be mapped to a function. Exactly the same mapping is used for the leaf PC and all the other PCs on the call stack. Most of the functions correspond directly to the source model of the program. Some do not; these functions are described in the following sections.

## Aliased Functions

Typically, functions are defined as global, meaning that their names are known everywhere in the program. The name of a global function must be unique within the executable. If there is more than one global function of a given name within the address space, the runtime linker resolves all references to one of them. The others are never executed, and so do not appear in the function list. In the Summary tab, you can see the shared object and object module that contain the selected function.

Under various circumstances, a function can be known by several different names. A very common example of this is the use of so-called weak and strong symbols for the same piece of code. A strong name is usually the same as the corresponding weak name, except that it has a leading underscore. Many of the functions in the threads library also have alternate names for pthreads and Solaris threads, as well as strong and weak names and alternate internal symbols. In all such cases, only one name is used in the function list of the Performance Analyzer. The name chosen is the last symbol at the given address in alphabetic order. This choice most often corresponds to the name that the user would use. In the Summary tab, all the aliases for the selected function are shown.

## Non-Unique Function Names

While aliased functions reflect multiple names for the same piece of code, there are circumstances under which multiple pieces of code have the same name:

- Sometimes, for reasons of modularity, functions are defined as static, meaning that their names are known only in some parts of the program (usually a single compiled object module). In such cases, several functions of the same name referring to quite different parts of the program appear in the Performance Analyzer. In the Summary tab, the object module name for each of these functions

is given to distinguish them from one another. In addition, any selection of one of these functions can be used to show the source, disassembly, and the callers and callees of that specific function.

- Sometimes a program uses wrapper or interposition functions that have the weak name of a function in a library and supersede calls to that library function. Some wrapper functions call the original function in the library, in which case both instances of the name appear in the Performance Analyzer function list. Such functions come from different shared objects and different object modules, and can be distinguished from each other in that way. The Collector wraps some library functions, and both the wrapper function and the real function can appear in the Performance Analyzer.

## Static Functions From Stripped Shared Libraries

Static functions are often used within libraries, so that the name used internally in a library does not conflict with a name that the user might use. When libraries are stripped, the names of static functions are deleted from the symbol table. In such cases, the Performance Analyzer generates an artificial name for each text region in the library containing stripped static functions. The name is of the form `<static>@0x12345`, where the string following the @ sign is the offset of the text region within the library. The Performance Analyzer cannot distinguish between contiguous stripped static functions and a single such function, so two or more such functions can appear with their metrics coalesced.

Stripped static functions are shown as called from the correct caller, except when the PC from the static function is a leaf PC that appears after the save instruction in the static function. Without the symbolic information, the Performance Analyzer does not know the save address, and cannot tell whether to use the return register as the caller. It always ignores the return register. Since several functions can be coalesced into a single `<static>@0x12345` function, the real caller or callee might not be distinguished from the adjacent functions.

## Fortran Alternate Entry Points

Fortran provides a way of having multiple entry points to a single piece of code, allowing a caller to call into the middle of a function. When such code is compiled, it consists of a prologue for the main entry point, a prologue to the alternate entry point, and the main body of code for the function. Each prologue sets up the stack for the function's eventual return and then branches or falls through to the main body of code.

The prologue code for each entry point always corresponds to a region of text that has the name of that entry point, but the code for the main body of the subroutine receives only one of the possible entry point names. The name received varies from one compiler to another.

The prologues rarely account for any significant amount of time, and the functions corresponding to entry points other than the one that is associated with the main body of the subroutine rarely appear in the Performance Analyzer. Call stacks representing time in Fortran subroutines with alternate entry points usually have PCs in the main body of the subroutine, rather than the prologue, and only the name associated with the main body will appear as a callee. Likewise, all calls from the subroutine are shown as being made from the name associated with the main body of the subroutine.

## Cloned Functions

The compilers have the ability to recognize calls to a function for which extra optimization can be performed. An example of such calls is a call to a function for which some of the arguments are constants. When the compiler identifies particular calls that it can optimize, it creates a copy of the function, which is called a clone, and generates optimized code. The clone function name is a mangled name that identifies the particular call. The Analyzer demangles the name, and presents each instance of a cloned function separately in the function list. Each cloned function has a different set of instructions, so the annotated disassembly listing shows the cloned functions separately. Each cloned function has the same source code, so the annotated source listing sums the data over all copies of the function.

## Inlined Functions

An inlined function is a function for which the instructions generated by the compiler are inserted at the call site of the function instead of an actual call. There are two kinds of inlining, both of which are done to improve performance, and both of which affect the Performance Analyzer.

- C++ inline function definitions. The rationale for inlining in this case is that the cost of calling a function is much greater than the work done by the inlined function, so it is better to simply insert the code for the function at the call site, instead of setting up a call. Typically, access functions are defined to be inlined, because they often only require one instruction. When you compile with the `-g` option, inlining of functions is disabled; compilation with `-g0` permits inlining of functions.

- Explicit or automatic inlining performed by the compiler at high optimization levels (4 and 5). Explicit and automatic inlining is performed even when `-g` is turned on. The rationale for this type of inlining can be to save the cost of a function call, but more often it is to provide more instructions for which register usage and instruction scheduling can be optimized.

Both kinds of inlining have the same effect on the display of metrics. Functions that appear in the source code but have been inlined do not show up in the function list, nor do they appear as callees of the functions into which they have been inlined. Metrics that would otherwise appear as inclusive metrics at the call site of the inlined function, representing time spent in the called function, are actually shown as exclusive metrics attributed to the call site, representing the instructions of the inlined function.

---

**Note** – Inlining can make data difficult to interpret, so you might want to disable inlining when you compile your program for performance analysis.

---

In some cases, even when a function is inlined, a so-called out-of-line function is left. Some call sites call the out-of-line function, but others have the instructions inlined. In such cases, the function appears in the function list but the metrics attributed to it represent only the out-of-line calls.

## Compiler-Generated Body Functions

When a compiler parallelizes a loop in a function, or a region that has parallelization directives, it creates new body functions that are not in the original source code. These functions are described in [“Parallel Execution and Compiler-Generated Body Functions” on page 124](#).

The Performance Analyzer shows these functions as normal functions, and assigns a name to them based on the function from which they were extracted, in addition to the compiler-generated name. Their exclusive and inclusive metrics represent the time spent in the body function. In addition, the function from which the construct was extracted shows inclusive metrics from each of the body functions. The means by which this is achieved is described in [“Parallel Execution Sequence” on page 125](#).

When a function containing parallel loops is inlined, the names of its compiler-generated body functions reflect the function into which it was inlined, not the original function.

---

**Note** – The names of compiler-generated body functions can only be demangled for modules compiled with `-g`

---

## Outline Functions

Outline functions can be created during feedback optimization. They represent code that is not normally executed. Specifically, it is code that is not executed during the training run that is used to generate the feedback. To improve paging and instruction-cache behavior, such code is moved elsewhere in the address space, and is made into a separate function. The name of the outline function encodes information about the section of outlined code, including the name of the function from which the code was extracted and the line number of the beginning of the section in the source code. These mangled names can vary from release to release. The Performance Analyzer provides a readable version of the function name.

Outline functions are not really called, but rather are jumped to; similarly they do not return, they jump back. In order to make the behavior more closely match the user's source code model, the Performance Analyzer imputes an artificial call from the main function to its outline portion.

Outline functions are shown as normal functions, with the appropriate inclusive and exclusive metrics. In addition, the metrics for the outline function are added as inclusive metrics in the function from which the code was outlined.

## Dynamically Compiled Functions

Dynamically compiled functions are functions that are compiled and linked while the program is executing. The Collector has no information about dynamically compiled functions that are written in C or C++, unless the user supplies the required information using the Collector API functions. See [“Dynamic Functions and Modules” on page 52](#) for information about the API functions. If information is not supplied, the function appears in the performance analysis tools as <Unknown>.

For Java programs, the Collector obtains information on methods that are compiled by the Java HotSpot™ virtual machine, and there is no need to use the API functions to provide the information. For other methods, the performance tools show information for the Java™ virtual machine that executes the methods. In Java mode, all methods are merged with the interpreted version. In Expert-Java mode, the HotSpot-compiled and interpreted versions of each method are both shown separately. In the machine mode, each HotSpot-compiled version is shown separately, and JVM functions are shown for each interpreted method.

## The <Unknown> Function

Under some circumstances, a PC does not map to a known function. In such cases, the PC is mapped to the special function named <Unknown>.

The following circumstances show PCs mapping to `<Unknown>`:

- When a function written in C or C++ is dynamically generated, and information about the function is not provided to the Collector using the Collector API functions. See [“Dynamic Functions and Modules” on page 52](#) for more information about the Collector API functions.
- When a Java method is dynamically compiled but Java profiling is disabled.
- When the PC corresponds to an address in the data section of the executable or a shared object. One case is the SPARC V7 version of `libc.so`, which has several functions (`.mul` and `.div`, for example) in its data section. The code is in the data section so that it can be dynamically rewritten to use machine instructions when the library detects that it is executing on a SPARC V8 or V9 platform.
- When the PC corresponds to a shared object in the address space of the executable that is not recorded in the experiment.
- When the PC is not within any known load object. The most likely cause of this is an unwind failure, where the value recorded as a PC is not a PC at all, but rather some other word. If the PC is the return register, and it does not seem to be within any known load object, it is ignored, rather than attributed to the `<Unknown>` function.
- When a PC maps to an internal part of the Java™ virtual machine for which the Collector has no symbolic information.

Callers and callees of the `<Unknown>` function represent the previous and next PCs in the call stack, and are treated normally.

## The `<no Java callstack recorded>` Function

The `<no Java callstack recorded>` function is similar to the `<Unknown>` function, but for Java threads, in the Java representation only. When the collector receives an event from a Java thread, it unwinds the native stack and calls into the Java virtual machine to obtain the corresponding Java stack. If that call fails for any reason, the event is shown in the analyzer with the artificial function `<no Java callstack recorded>`. The JVM may refuse to report a callstack either to avoid deadlock, or when unwinding the Java stack would cause excessive synchronization.

## The `<Total>` Function

The `<Total>` function is an artificial construct used to represent the program as a whole. All performance metrics, in addition to being attributed to the functions on the call stack, are attributed to the special function `<Total>`. It appears at the top of the function list and its data can be used to give perspective on the data for other functions. In the Callers-Callees list, it is shown as the nominal caller of `_start()`



in the main thread of execution of any program, and also as the nominal caller of `_thread_start()` for created threads. If the stack unwind was incomplete, the `<Total>` function can appear as the caller of other functions.

### Functions Related to HW Counter Profiling

The following functions are related to HW counter profiling:

- `collector_not_program_related`: The counter does not relate to the program.
- `collector_lost_hwc_overflow`: The counter appears to have exceeded the overflow value without generating an overflow signal. The value is recorded and the counter reset.
- `collector_lost_sigemt`: The counter appears to have exceeded the overflow value and been halted but the overflow signal appears to be lost. The value is recorded and the counter reset.
- `collector_hwc_ABORT`: Reading the hardware counters has failed, typically when a privileged process has taken control of the counters, resulting in the termination of hardware counter collection.
- `collector_final_counters`: The values of the counters taken immediately before suspending or terminating collection, with the count since the previous overflow. If this corresponds to a significant fraction of the `<Total>` count, a smaller overflow interval (that is, a higher resolution configuration) is recommended.
- `collector_record_counters`: The counts accumulated while handling and recording hardware counter events, partially accounting for hardware counter profiling overhead. If this corresponds to a significant fraction of the `<Total>` count, a larger overflow interval (that is, a lower resolution configuration) is recommended.

---

## Mapping Data Addresses to Program Data Objects

Once a PC from a hardware counter event corresponding to a memory operation has been processed to successfully backtrack to a likely causal memory-referencing instruction, the Performance Analyzer uses instruction identifiers and descriptors provided by the compiler in its hardware profiling support information to derive the associated program data object.

The term *dataobject* is used to refer to program constants, variables, arrays and aggregates such as structures and unions, along with distinct aggregate elements, described in source code. Depending on the source language, dataobject types and

their sizes will vary. Many dataobjects are explicitly named in source programs, while others may be unnamed. Some dataobjects are derived or aggregated from other (simpler) dataobjects, resulting in a rich, often complex, set of dataobjects.

Each dataobject has an associated scope, the region of the source program where it is defined and can be referenced, which may be global (such as a load object), a particular compilation unit (an object file), or function. Identical data objects may be defined with different scopes, or particular data objects referred to differently in different scopes.

Data-derived metrics from hardware counter events for memory operations collected with backtracking enabled are attributed to the associated program dataobject and propagate to any aggregates containing the dataobject and the artificial `<Total>` which is considered to contain all dataobjects (including `<Unknown>` and `<Scalars>`). The different subtypes of `<Unknown>` propagate up to the `<Unknown>` aggregate. The following section describes the `<Total>`, `<Scalars>`, and `<Unknown>` dataobjects.

## Dataobject Descriptors

Data objects are fully described by a combination of their declared type and name. A simple scalar dataobject `{int i}` describes an variable called `i` of type `int`, while `{const+pointer+int p}` describes a constant pointer to a type `int` called `p`. Types containing spaces have them replaced with underscore (`_`), and unnamed data objects are represented with a name of dash (`-`), for example: `{double_precision_complex -}`.

An entire aggregate is similarly represented `{structure:foo_t foo}` for a structure of type `foo_t` called `foo`. A single element of an aggregate requires the additional specification of its container, for example, `{structure:foo_t foo}. {int i}` for a member `i` of type `int` of the previous structure `foo`. Aggregates may also themselves be elements of (larger) aggregates, with their corresponding descriptor constructed as a concatenation of aggregate descriptors and ultimately a scalar descriptor.

While a fully-qualified descriptor may not always be necessary to disambiguate dataobjects, it provides a generic complete specification to assist with dataobject identification.

## The <Total> Dataobject

The <Total> dataobject is an artificial construct used to represent the program's data objects as a whole. All performance metrics, in addition to being attributed to a distinct dataobject (and any aggregate to which it belongs), are attributed to the special dataobject <Total>. It appears at the top of the dataobject list and its data can be used to give perspective to the data for other data objects.

## The <Scalars> Dataobject

While aggregate elements have their performance metrics additionally attributed into the metric value for their associated aggregate, all of the scalar constants and variables have their performance metrics additionally attributed into the metric value for the artificial <Scalars> dataobject.

## The <Unknown> Dataobject and Its Elements

Under various circumstances, event data can not be mapped to a particular dataobject. In such cases, the data is mapped to the special dataobject named <Unknown> and one of its elements as described below.

- Unascertainable

One or more compilation objects were compiled without hardware profiling support, such that it is not possible to ascertain data objects associated with memory-referencing instructions or validate backtracking.

- Unverifiable

The hardware profiling support information provided in the compilation object was insufficient to verify the validity of backtracking.

- Unresolvable

Event backtracking encountered a control transfer target and, being unable to resolve whether the control transfer actually occurred or not, was unable to determine the likely causal memory-referencing instruction (and its associated dataobject).

- Unspecified

Backtracking determined the likely causal memory-referencing instruction, but its associated dataobject was not specified by the compiler.

- Unidentified

Backtracking determined the likely causal memory-referencing instruction, but it was not identified by the compiler and associated dataobject determination is therefore not possible. Compiler temporaries are generally unidentified.

---

# Annotated Code Listings

Annotated source code and annotated disassembly code are useful for determining which source lines or instructions within a function are responsible for poor performance. This section describes the annotation process and some of the issues involved in interpreting the annotated code.

## Annotated Source Code

Annotated source code shows the resource consumption of an application at the source-line level. It is produced by taking the PCs that are recorded in the application's call stack, and mapping each PC to a source line. To produce an annotated source file, the Performance Analyzer first determines all of the functions that are generated in a particular object module (.o file) or load object, then scans the data for all PCs from each function. In order to produce annotated source, the Performance Analyzer must be able to find and read the object module or load object to determine the mapping from PCs to source lines, and it must be able to read the source file to produce an annotated copy, which is displayed. The Performance Analyzer searches for the source, object and executable files in the following locations in turn, and stops when it finds a file of the correct basename:

- The archive directories of experiments.
- The current working directory.
- The absolute pathname as recorded in the executables or compilation objects.

The compilation process goes through many stages, depending on the level of optimization requested, and transformations take place which can confuse the mapping of instructions to source lines. For some optimizations, source line information might be completely lost, while for others, it might be confusing. The compiler relies on various heuristics to track the source line for an instruction, and these heuristics are not infallible.

## Interpreting Source Line Metrics

Metrics for an instruction must be interpreted as metrics accrued while waiting for the instruction to be executed. If the instruction being executed when an event is recorded comes from the same source line as the leaf PC, the metrics can be interpreted as due to execution of that source line. However, if the leaf PC comes from a different source line than the instruction being executed, at least some of the metrics for the source line that the leaf PC belongs to must be interpreted as metrics accumulated while this line was waiting to be executed. An example is when a value that is computed on one source line is used on the next source line.

The issue of how to interpret the metrics matters most when there is a substantial delay in execution, such as at a cache miss or a resource queue stall, or when an instruction is waiting for a result from a previous instruction. In such cases the metrics for the source lines can seem to be unreasonably high, and you should look at other lines in the code to find the line responsible for the high metric value.

## Metric Formats

The four possible formats for the metrics that can appear on a line of annotated source code are explained in [TABLE 5-3](#).

**TABLE 5-3** Annotated Source-Code Metrics

Metric	Significance
(Blank)	No PC in the program corresponds to this line of code. This case should always apply to comment lines, and applies to apparent code lines in the following circumstances: <ul style="list-style-type: none"><li>• All the instructions from the apparent piece of code have been eliminated during optimization.</li><li>• The code is repeated elsewhere, and the compiler performed common subexpression recognition and tagged all the instructions with the lines for the other copy.</li><li>• The compiler tagged an instruction with an incorrect line number.</li></ul>
0.	Some PCs in the program were tagged as derived from this line, but there was no data that referred to those PCs: they were never in a call stack that was sampled statistically or traced for thread-synchronization data. The 0. metric does not indicate that the line was not executed, only that it did not show up statistically in a profiling data packet or a tracing data packet.
0.000	At least one PC from this line appeared in the data, but the computed metric value rounded to zero.
1.234	The metrics for all PCs attributed to this line added up to the non-zero numerical value shown.

## Compiler Commentary

Various parts of the compiler can incorporate commentary into the executable. Each comment is associated with a specific line of source code. When the annotated source is written, the compiler commentary for any source line appears immediately preceding the source line.

The compiler commentary describes many of the transformations which have been made to the source code to optimize it. These transformations include loop optimizations, parallelization, inlining and pipelining.

## Common Subexpression Elimination

One very common optimization recognizes that the same expression appears in more than one place, and that performance can be improved by generating the code for that expression in one place. For example, if the same operation appears in both the `if` and the `else` branches of a block of code, the compiler can move that operation to just before the `if` statement. When it does so, it assigns line numbers to the instructions based on one of the previous occurrences of the expression. If the line numbers assigned to the common code correspond to one branch of an `if` structure, and the code actually always takes the other branch, the annotated source shows metrics on lines within the branch that is not taken.

## Parallelization Directives

When the compiler generates body functions from code that contains parallelization directives, inclusive metrics for the parallel loop or section are attributed to the parallelization directive, because this line is the call site for the compiler-generated body function. Inclusive and exclusive metrics also appear on the code in the loops or sections. These metrics sum to the inclusive metrics on the parallelization directives.

## Special Lines in the Source

Whenever the source line for a PC cannot be determined, the metrics for that PC are attributed to a special source line that is inserted at the top of the annotated source file. High metrics on that line indicates that part of the code from the given object module does not have line-mappings. Annotated disassembly can help you determine the instructions that do not have mappings. The special lines are:

- *Function* <instructions without line numbers>  
Where *function* is the name of the function from which the instruction came.
- *Function* <source file name not recorded>  
Where *function* is the name of the function from which the instruction came.

## Annotated Disassembly Code

Annotated disassembly provides an assembly-code listing of the instructions of a function or object module, with the performance metrics associated with each instruction. Annotated disassembly can be displayed in several ways, determined by whether line-number mappings and the source file are available, and whether the object module for the function whose annotated disassembly is being requested is known:

- If the object module is not known, the Performance Analyzer disassembles the instructions for just the specified function, and does not show any source lines in the disassembly.
- If the object module is known, the disassembly covers all functions within the object module.
- If the source file is available, and line number data is recorded, the Performance Analyzer can interleave the source with the disassembly, depending on the display preference.
- If the compiler has inserted any commentary into the object code, it too, is interleaved in the disassembly if the corresponding preferences are set.

Each instruction in the disassembly code is annotated with the following information.

- A source line number, as reported by the compiler
- Its relative address
- The hexadecimal representation of the instruction, if requested
- The assembler ASCII representation of the instruction

Where possible, call addresses are resolved to symbols (such as function names). Metrics are shown on the lines for instructions, and can be shown on any interleaved source code if the corresponding preference is set. Possible metric values are as described for source-code annotations, in [TABLE 5-3](#).

When code is not optimized, the line numbers for each instruction are in sequential order, and the interleaving of source lines and disassembled instructions occurs in the expected way. When optimization takes place, instructions from later lines sometimes appear before those from earlier lines. The Performance Analyzer's algorithm for interleaving is that whenever an instruction is shown as coming from line *N*, all source lines up to and including line *N* are written before the instruction. One effect of optimization is that source code can appear between a control transfer instruction and its delay slot instruction. Compiler commentary associated with line *N* of the source is written immediately before that line.

Interpreting annotated disassembly is not straightforward. The leaf PC is the address of the next instruction to execute, so metrics attributed to an instruction should be considered as time spent waiting for the instruction to execute. However, the execution of instructions does not always happen in sequence, and there might be delays in the recording of the call stack. To make use of annotated disassembly, you should become familiar with the hardware on which you record your experiments and the way in which it loads and executes instructions.

The next few subsections discuss some of the issues of interpreting annotated disassembly.

## Instruction Issue Grouping

Instructions are loaded and issued in groups known as instruction issue groups. Which instructions are in the group depends on the hardware, the instruction type, the instructions already being executed, and any dependencies on other instructions or registers. This means that some instructions might be underrepresented because they are always issued in the same clock cycle as the previous instruction, so they never represent the next instruction to be executed. It also means that when the call stack is recorded, there might be several instructions which could be considered the next instruction to execute.

Instruction issue rules vary from one processor type to another, and depend on the instruction alignment within cache lines. Since the linker forces instruction alignment at a finer granularity than the cache line, changes in a function that might seem unrelated can cause different alignment of instructions. The different alignment can cause a performance improvement or degradation.

The following artificial situation shows the same function compiled and linked in slightly different circumstances. The two output examples shown below are the annotated disassembly listings from `er_print`. The instructions for the two examples are identical, but the instructions are aligned differently.



In this example the instruction alignment maps the two instructions `cmp` and `bl`, a to different cache lines, and a significant amount of time is used waiting to execute these two instructions.

Excl. User CPU sec.	Incl. User CPU sec.	
		1. static int
		2. ifunc()
		3. {
		4.     int i;
		5.
		6.     for (i=0; i<10000; i++)
		<function: ifunc>
0.010	0.010	[ 6]   1066c:  clr            %o0
0.	0.	[ 6]   10670:  sethi        %hi(0x2400), %o5
0.	0.	[ 6]   10674:  inc            784, %o5
		7.     i++;
0.	0.	[ 7]   10678:  inc            2, %o0
## 1.360	1.360	[ 7]   1067c:  cmp            %o0, %o5
## 1.510	1.510	[ 7]   10680:  bl,a          0x1067c
0.	0.	[ 7]   10684:  inc            2, %o0
0.	0.	[ 7]   10688:  retl
0.	0.	[ 7]   1068c:  nop
		8.     return i;
		9. }

In this example, the instruction alignment maps the two instructions `cmp` and `bl,a` to the same cache line, and a significant amount of time is used waiting to execute only one of these instructions.

Excl. User CPU sec.	Incl. User CPU sec.	
		1. static int
		2. ifunc()
		3. {
		4.     int i;
		5.
		6.     for (i=0; i<10000; i++)
		<function: ifunc>
0.	0.	[ 6] 10684: clr            %o0
0.	0.	[ 6] 10688: sethi         %hi(0x2400), %o5
0.	0.	[ 6] 1068c: inc           784, %o5
		7.     i++;
0.	0.	[ 7] 10690: inc            2, %o0
## 1.440	1.440	[ 7] 10694: cmp          %o0, %o5
0.	0.	[ 7] 10698: bl,a         0x10694
0.	0.	[ 7] 1069c: inc           2, %o0
0.	0.	[ 7] 106a0: retl
0.	0.	[ 7] 106a4: nop
		8.     return i;
		9. }

## Instruction Issue Delay

Sometimes, specific leaf PCs appear more frequently because the instruction that they represent is delayed before issue. This can occur for a number of reasons, some of which are listed below:

- The previous instruction takes a long time to execute and is not interruptible, for example when an instruction traps into the kernel.
- An arithmetic instruction needs a register that is not available because the register contents were set by an earlier instruction that has not yet completed. An example of this sort of delay is a load instruction that has a data cache miss.
- A floating-point arithmetic instruction is waiting for another floating-point instruction to complete. This situation occurs for instructions that cannot be pipelined, such as square root and floating-point divide.
- The instruction cache does not include the memory word that contains the instruction (I-cache miss).

- On UltraSPARC III processors, a cache miss on a load instruction blocks all instructions that follow it until the miss is resolved, regardless of whether these instructions use the data item that is being loaded. UltraSPARC® II processors only block instructions that use the data item that is being loaded.

## Attribution of Hardware Counter Overflows

Apart from TLB misses, the call stack for a hardware counter overflow event is recorded at some point further on in the sequence of instructions than the point at which the overflow occurred, for various reasons including the time taken to handle the interrupt generated by the overflow. For some counters, such as cycles or instructions issued, this does not matter. For other counters, such as those counting cache misses or floating point operations, the metric is attributed to a different instruction from that which is responsible for the overflow. Often the PC that caused the event is only a few instructions before the recorded PC, and the instruction can be correctly located in the disassembly listing. However, if there is a branch target within this instruction range, it might be difficult or impossible to tell which instruction corresponds to the PC that caused the event. For hardware counters that count memory access events, the Collector searches for the PC that caused the event if the counter name is prefixed with a plus, +.

### Special Lines in the Disassembly and PCs Tabs

There are a number of special lines that may appear in the Disassembly and PCs tabs. This section describes such lines using examples as they might appear in the analyzer.

- `<static>@0x1a4400 + 0x000032B8`  
Represents a static function and its offset
- `<library.so> -- no functions found + 0x0000F870`  
As above, with offset
- `Method_Name <HotSpot-compiled leaf instructions>`  
Represents the HotSpot-compiled version of the named Java method
- `Method_Name <Java native method>`  
Represents instructions generated by a Java native method
- `<no Java callstack recorded> + 0x00000000`  
The offset is an error code from the JVM, usually 0 indicating nothing -- but no unwind
- `<branch target>`  
In disassembly, where the backtracking for event PC and its effective address fails because it runs into a branch target, and so can not back up past it.

## Program Linkage Table (PLT) Instructions

When a function in one load object calls a function in a different shared object, the actual call transfers first to a three-instruction sequence in the PLT, and then to the real destination. The analyzer removes PCs that correspond to the PLT, and assigns the metrics for these PCs to the call instruction. Therefore, if a call instruction has an unexpectedly high metric, it could be due to the PLT instructions rather than the call instructions. See also [“Function Calls Between Shared Objects”](#) on page 118.

# Manipulating Experiments and Viewing Annotated Code Listings

---

This chapter describes the utilities which are available for use with the Collector and Performance Analyzer.

This chapter covers the following topics:

- [Manipulating Experiments](#)
  - [Viewing Annotated Code Listings With `er\_src`](#)
  - [Other Utilities](#)
- 

## Manipulating Experiments

Experiments are stored in a directory that is created by the Collector. To manipulate experiments, you can use the usual Unix commands `cp`, `mv` and `rm` and apply them to the directory. This is not true for experiments from releases earlier than Forte Developer 7 (Sun™ ONE Studio 7, Enterprise Edition for the Solaris Operating System). Three utilities which behave like the Unix commands have been provided to copy, move and delete experiments. These utilities are `er_cp(1)`, `er_mv(1)` and `er_rm(1)`, and are described below.

The data in the experiment includes archive files for each of the load objects used by your program. These archive files contain the absolute path of the load object and the date on which it was last modified. This information is not changed when you move or copy an experiment.

```
er_cp [-V] experiment1 experiment2
```

```
er_cp [-V] experiment-list directory
```

The first form of the `er_cp` command copies *experiment1* to *experiment2*. If *experiment2* exists, `er_cp` exits with an error message. The second form copies a blank-separated list of experiments to a directory. If the directory already contains an experiment with the same name as one of the experiments being copied, `er_mv` exits with an error message. The `-v` option prints the version of `er_cp`. This command does not copy experiments created with software releases earlier than the Forte Developer 7 release.

```
er_mv [-V] experiment1 experiment2
```

```
er_mv [-V] experiment-list directory
```

The first form of the `er_mv` command moves *experiment1* to *experiment2*. If *experiment2* exists, `er_mv` exits with an error message. The second form moves a blank-separated list of experiments to a directory. If the directory already contains an experiment with the same name as one of the experiments being moved, `er_mv` exits with an error message. The `-v` option prints the version of `er_mv`. This command does not move experiments created with software releases earlier than the Forte Developer 7 release.

```
er_rm [-f] [-V] experiment-list
```

Removes a list of experiments or experiment groups. When experiment groups are removed, each experiment in the group is removed then the group file is removed. The `-f` option suppresses error messages and ensures successful completion, whether or not the experiments are found. The `-v` option prints the version of `er_rm`. This command removes experiments created with software releases earlier than the Forte Developer 7 release.

---

## Viewing Annotated Code Listings With `er_src`

Annotated source code and annotated disassembly code can be viewed using the `er_src` utility, without running an experiment. The display is generated in the same way as in the Performance Analyzer, except that it does not display any metrics. The syntax of the `er_src` command is

```
er_src [ options ] object item tag
```

*object* is the name of an executable, a shared object, or an object file (.o file).

*item* is the name of a function or of a source or object file used to build the executable or shared object; it can be omitted when an object file is specified.

*tag* is an index used to determine which *item* is being referred to when multiple functions have the same name. If it is not needed, it can be omitted. If it is needed and is omitted, a message listing the possible choices is printed.

The following sections describe the options accepted by the `er_src` utility.

### `-C` *commentary-classes*

Define the compiler commentary classes to be shown. *commentary-classes* is a list of classes separated by colons. See [“Commands That Control the Source and Disassembly Listings” on page 92](#) for a description of these classes.

The commentary classes can be specified in a defaults file. The system wide `er.rc` defaults file is read first, then a `.er.rc` file in the user’s home directory, if present, then a `.er.rc` file in the current directory. Defaults from the `.er.rc` file in your home directory override the system defaults, and defaults from the `.er.rc` file in the current directory override both home and system defaults. These files are also used by the Performance Analyzer and `er_print`, but only the settings for source and disassembly compiler commentary are used by `er_src`.

See [“Commands That Set Defaults” on page 102](#) for a description of the defaults files. Commands in a defaults file other than `scc` and `dcc` are ignored by `er_src`.

### `-d`

Include the disassembly in the listing. The default listing does not include the disassembly. If there is no source available, a listing of the disassembly without compiler commentary is produced.

### `-o` *filename*

Open the file *filename* for output of the listing. By default, output is written to `stdout`.

### `-V`

Print the current release version.

---

## Other Utilities

There are some other utilities that should not need to be used in normal circumstances. They are documented here for completeness, with a description of the circumstances in which it might be necessary to use them.

### The `er_archive` Utility

The syntax of the `er_archive` command is as follows.

```
er_archive [-qAF] experiment
er_archive -V
```

The `er_archive` utility is automatically run when an experiment completes normally, or when the Performance Analyzer or `er_print` command is started on an experiment. It reads the list of shared objects referenced in the experiment, and constructs an archive file for each. Each output file is named with a suffix of `.archive`, and contains function and module mappings for the shared object.

If the target program terminates abnormally, `er_archive` might not be run by the Collector. If you want to examine the experiment from an abnormally-terminated run on a different machine from the one on which it was recorded, you must run `er_archive` on the experiment, on the machine on which the data was recorded. To ensure that the load objects are available on the machine to which the experiment is copied, use the `-A` option.

An archive file is generated for all shared objects referred to in the experiment. These archives contain the addresses, sizes and names of each object file and each function in the load object, as well as the absolute path of the load object and a time stamp for its last modification.

If the shared object cannot be found when `er_archive` is run, or if it has a time stamp differing from that recorded in the experiment, or if `er_archive` is run on a different machine from that on which the experiment was recorded, the archive file contains a warning. Warnings are also written to `stderr` whenever `er_archive` is run manually (without the `-q` flag).

The following sections describe the options accepted by the `er_archive` utility.



-Q

Do not write any warnings to `stderr`. Warnings are incorporated into the archive file, and shown in the Performance Analyzer or `er_print` output.

-A

Request writing of all load objects into the experiment. This argument can be used to generate experiments that are more readily copied to a machine other than the one on which the experiment was recorded.

-F

Force writing or rewriting of archive files. This argument can be used to run `er_archive` by hand, to rewrite files that had warnings.

-V

Write version number information for `er_archive` and exit.

## The `er_export` Utility

The syntax of the `er_export` command is as follows.

```
er_export [-V] experiment
```

The `er_export` utility converts the raw data in an experiment into ASCII text. The format and the content of the file are subject to change, and should not be relied on for any use. This utility is intended to be used only when the Performance Analyzer cannot read an experiment; the output allows the tool developers to understand the raw data and analyze the failure. The `-V` option prints version number information.



## Profiling Programs With `prof`, `gprof`, and `tcov`

---

The tools discussed in this appendix are standard utilities for timing programs and obtaining performance data to analyze, and are called traditional profiling tools. The profiling tools `prof` and `gprof` are provided with the Solaris Operating System. `tcov` is a code coverage tool provided with the Sun compilers and tools.

---

**Note** – If you want to track how many times a function is called or how often a line of source code is executed, use the traditional profiling tools. If you want a detailed analysis of where your program is spending time, you can get more accurate information using the Collector and Performance Analyzer. See [Chapter 3](#) and the online help for information on using these tools.

---

[TABLE A-1](#) describes the information that is generated by these standard performance profiling tools.

**TABLE A-1** Performance Profiling Tools

---

Command	Output
<code>prof</code>	Generates a statistical profile of the CPU time used by a program and an exact count of the number of times each function is entered.
<code>gprof</code>	Generates a statistical profile of the CPU time used by a program, along with an exact count of the number of times each function is entered and the number of times each arc (caller-callee pair) in the program's call graph is traversed.
<code>tcov</code>	Generates exact counts of the number of times each statement in a program is executed.

---

Not all the traditional profiling tools work on modules written in programming languages other than C. See the sections on each tool for more information about languages.

This appendix covers the following topics:

- [Using `prof` to Generate a Program Profile](#)
- [Using `gprof` to Generate a Call Graph Profile](#)
- [Using `tcov` for Statement-Level Analysis](#)
- [Using `tcov` Enhanced for Statement-Level Analysis](#)

---

## Using `prof` to Generate a Program Profile

`prof` generates a statistical profile of the CPU time used by a program and counts the number of times each function in a program is entered. Different or more detailed data is provided by the `gprof` call-graph profile and the `tcov` code coverage tools.

To generate a profile report using `prof`:

**1. Compile your program with the `-p` compiler option.**

**2. Run your program.**

Profiling data is sent to a profile file called `mon.out`. This file is overwritten each time you run the program.

**3. Run `prof` to generate a profile report.**

The syntax of the `prof` command is as follows.

```
% prof program-name
```

Here, *program-name* is the name of the executable. The profile report is written to `stdout`. It is presented as a series of rows for each function under these column headings:

- `%Time`—The percentage of the total CPU time consumed by this function.
- `Seconds`—The total CPU time accounted for by this function.
- `Cumsecs`—A running sum of the number of seconds accounted for by this function and those listed before it.
- `#Calls`—The number of times this function is called.
- `msecs/call`—The average number of milliseconds this function consumes each time it is called.
- `Name`—The name of the function.

The use of `prof` is illustrated in the following example.

```
% cc -p -o index.assist index.assist.c
% index.assist
% prof index.assist
```

The profile report from `prof` is shown in the table below:

%Time	Seconds	Cumsecs	#Calls	msecs/call	Name
19.4	3.28	3.28	11962	0.27	<code>compare_strings</code>
15.6	2.64	5.92	32731	0.08	<code>_strlen</code>
12.6	2.14	8.06	4579	0.47	<code>__doprnt</code>
10.5	1.78	9.84			<code>mcount</code>
9.9	1.68	11.52	6849	0.25	<code>_get_field</code>
5.3	0.90	12.42	762	1.18	<code>_fgets</code>
4.7	0.80	13.22	19715	0.04	<code>_strcmp</code>
4.0	0.67	13.89	5329	0.13	<code>_malloc</code>
3.4	0.57	14.46	11152	0.05	<code>_insert_index_entry</code>
3.1	0.53	14.99	11152	0.05	<code>_compare_entry</code>
2.5	0.42	15.41	1289	0.33	<code>lmodt</code>
0.9	0.16	15.57	761	0.21	<code>_get_index_terms</code>
0.9	0.16	15.73	3805	0.04	<code>_strcpy</code>
0.8	0.14	15.87	6849	0.02	<code>_skip_space</code>
0.7	0.12	15.99	13	9.23	<code>_read</code>
0.7	0.12	16.11	1289	0.09	<code>ldivt</code>
0.6	0.10	16.21	1405	0.07	<code>_print_index</code>
.					
.					
.					

(The rest of the output is insignificant)

The profile report shows that most of the program execution time is spent in the `compare_strings()` function; after that, most of the CPU time is spent in the `_strlen()` library function. To make this program more efficient, the user would concentrate on the `compare_strings()` function, which consumes nearly 20% of the total CPU time, and improve the algorithm or reduce the number of calls.

It is not obvious from the `prof` profile report that `compare_strings()` is heavily recursive, but you can deduce this by using the call graph profile described in “Using `gprof` to Generate a Call Graph Profile” on page 158. In this particular case, improving the algorithm also reduces the number of calls.

---

**Note** – For the Solaris Operating System versions 7 and 8, the profile of CPU time is accurate for programs that use multiple CPUs, but the fact that the counts are not locked may affect the accuracy of the counts for functions.

---

---

## Using `gprof` to Generate a Call Graph Profile

While the flat profile from `prof` can provide valuable data for performance improvements, a more detailed analysis can be obtained by using a call graph profile to display a list identifying which modules are called by other modules, and which modules call other modules. Sometimes removing calls altogether can result in performance improvements.

---

**Note** – `gprof` attributes the time spent within a function to the callers in proportion to the number of times that each arc is traversed. Because all calls are not equivalent in performance, this behavior might lead to incorrect assumptions. See the performance analyzer tutorial on `developers.sun.com` for an example.

---

Like `prof`, `gprof` generates a statistical profile of the CPU time that is used by a program and it counts the number of times that each function is entered. `gprof` also counts the number of times that each arc in the program’s call graph is traversed. An *arc* is a caller-callee pair.

---

**Note** – The profile of CPU time is accurate for programs that use multiple CPUs for the Solaris Operating System versions 7 and 8, but the fact that the counts are not locked may affect the accuracy of the counts for functions.

---

To generate a profile report using `gprof`:

1. **Compile your program with the appropriate compiler option.**
  - For C programs, use the `-xpg` option.
  - For Fortran programs, use the `-pg` option.

## 2. Run your program.

Profiling data is sent to a profile file called `gmon.out`. This file is overwritten each time you run the program.

## 3. Run `gprof` to generate a profile report.

The syntax of the `prof` command is as follows.

```
% gprof program-name
```

Here, *program-name* is the name of the executable. The profile report is written to `stdout`, and can be large. The report consists of two major items:

- The full call graph profile, which shows information about the callers and callees of each function in the program. The format is illustrated in the example given below.
- The flat profile, which is similar to the summary the `prof` command supplies.

The profile report from `gprof` contains an explanation of what the various parts of the summary mean and identifies the granularity of the sampling, as shown in the following example.

```
granularity: each sample hit covers 4 byte(s) for 0.07% of 14.74
seconds
```

The 4 bytes means resolution to a single instruction. The 0.07% of 14.74 seconds means that each sample, representing ten milliseconds of CPU time, accounts for 0.07% of the run.

The use of `gprof` is illustrated in the following example.

```
% cc -xpg -o index.assist index.assist.c
% index.assist
% gprof index.assist > g.output
```

The following table is part of the call graph profile.

index	%time	self	descendants	called/total parents	name	index
				called/total children		
-----						

		0.00	14.47	1/1	start	[1]
[2]	98.2	0.00	14.47	1	_main	[2]
		0.59	5.70	760/760	_insert_index_entry	[3]
		0.02	3.16	1/1	_print_index	[6]
		0.20	1.91	761/761	_get_index_terms	[11]
		0.94	0.06	762/762	_fgets	[13]
		0.06	0.62	761/761	_get_page_number	[18]
		0.10	0.46	761/761	_get_page_type	[22]
		0.09	0.23	761/761	_skip_start	[24]
		0.04	0.23	761/761	_get_index_type	[26]
		0.07	0.00	761/820	_insert_page_entry	[34]
-----						
				10392	_insert_index_entry	[3]
		0.59	5.70	760/760	_main	[2]
[3]	42.6	0.59	5.70	760+10392	_insert_index_entry	[3]
		0.53	5.13	11152/11152	_compare_entry	[4]
		0.02	0.01	59/112	_free	[38]
		0.00	0.00	59/820	_insert_page_entry	[34]
				10392	_insert_index_entry	[3]
-----						

In this example there are 761 lines of data in the input file to the `index.assist` program. The following conclusions can be made:

- `fgets()` is called 762 times. The last call to `fgets()` returns an end-of-file.
- The `insert_index_entry()` function is called 760 times from `main()`.
- In addition to the 760 times that `insert_index_entry()` is called from `main()`, `insert_index_entry()` also calls itself 10,392 times. `insert_index_entry()` is heavily recursive.
- `compare_entry()`, which is called from `insert_index_entry()`, is called 11,152 times, which is equal to 760+10,392 times. There is one call to `compare_entry()` for every time that `insert_index_entry()` is called. This is correct. If there were a discrepancy in the number of calls, you would suspect some problem in the program logic.



- `insert_page_entry()` is called 820 times in total: 761 times from `main()` while the program is building index nodes, and 59 times from `insert_index_entry()`. This frequency indicates that there are 59 duplicated index entries, so their page number entries are linked into a chain with the index nodes. The duplicate index entries are then freed; hence the 59 calls to `free()`.

---

## Using `tcov` for Statement-Level Analysis

The `tcov` utility gives information on how often a program executes segments of code. It produces a copy of the source file, annotated with execution frequencies. The code can be annotated at the basic block level or the source line level. A basic block is a linear segment of source code with no branches. The statements in a basic block are executed the same number of times, so a count of basic block executions also tells you how many times each statement in the block was executed. The `tcov` utility does not produce any time-based data.

---

**Note** – Although `tcov` works with both C and C++ programs, it does not support files that contain `#line` or `#file` directives. `tcov` does not enable test coverage analysis of the code in the `#include` header files.

---

To generate annotated source code using `tcov`:

### 1. Compile your program with the appropriate compiler option.

- For C programs, use the `-xa` option.
- For Fortran and C++ programs, use the `-a` option.

If you compile with the `-a` or `-xa` option you must also link with it. The compiler creates a coverage data file with the suffix `.d` for each object file. The coverage data file is created in the directory specified by the environment variable `TCOVDIR`. If `TCOVDIR` is not set, the coverage data file is created in the current directory.

---

**Note** – Programs compiled with `-xa` (C) or `-a` (other compilers) run more slowly than they normally would, because updating the `.d` file for each execution takes considerable time.

---

### 2. Run your program.

When your program completes, the coverage data files are updated.

### 3. Run `tcov` to generate annotated source code.

The syntax of the `tcov` command is as follows.

```
% tcov options source-file-list
```

Here, *source-file-list* is a list of the source code filenames. For a list of options, see the `tcov(1)` man page. The default output of `tcov` is a set of files, each with the suffix `.tcov`, which can be changed with the `-o filename` option.

A program compiled for code coverage analysis can be run multiple times (with potentially varying input); `tcov` can be used on the program after each run to compare behavior.

The following example illustrates the use of `tcov`.

```
% cc -xa -o index.assist index.assist.c
% index.assist
% tcov index.assist.c
```

This small fragment of the C code from one of the modules of `index.assist` shows the `insert_index_entry()` function, which is called recursively. The numbers to the left of the C code show how many times each basic block was executed. The `insert_index_entry()` function is called 11,152 times.

```

11152      struct index_entry *
-> insert_index_entry(node, entry)
      struct index_entry *node;
      struct index_entry *entry;
      {
          int result;
          int level;

          result = compare_entry(node, entry);
          if (result == 0) {      /* exact match */
                                  /* Place the page entry for the
duplicate */
                                  /* into the list of pages for this node
*/
59          ->          insert_page_entry(node, entry->page_entry);
                          free(entry);
                          return(node);
          }

11093      ->          if (result > 0)      /* node greater than new entry -- */
                                  /* move to lesser nodes */
3956      ->          if (node->lesser != NULL)
3626      ->          insert_index_entry(node->lesser, entry);
          else {
330      ->          node->lesser = entry;
                          return (node->lesser);
          }
          else      /* node less than new entry -- */
                                  /* move to greater nodes */
7137      ->          if (node->greater != NULL)
6766      ->          insert_index_entry(node->greater, entry);
          else {
371      ->          node->greater = entry;
                          return (node->greater);
          }
      }

```

The `tcov` utility places a summary like the following at the end of the annotated program listing. The statistics for the most frequently executed basic blocks are listed in order of execution frequency. The line number is the number of the first line in the block.

The following is the summary for the `index.assist` program:

Top 10 Blocks

Line	Count
240	21563
241	21563
245	21563
251	21563
250	21400
244	21299
255	20612
257	16805
123	12021
124	11962

77        Basic blocks in this file  
55        Basic blocks executed  
71.43    Percent of the file executed

439144    Total basic block executions  
5703.17   Average executions per basic block

## Creating `tcov` Profiled Shared Libraries

It is possible to create a `tcov` profiled shareable library and use it in place of the corresponding library in binaries which have already been linked. Include the `-xa` (C) or `-a` (other compilers) option when creating the shareable libraries, as shown in this example.

```
% cc -G -xa -o foo.so.1 foo.o
```

This command includes a copy of the `tcov` profiling functions in the shareable libraries, so that clients of the library do not need to relink. If a client of the library is already linked for profiling, then the version of the `tcov` functions used by the client is used to profile the shareable library.

## Locking Files

`tcov` uses a simple file-locking mechanism for updating the block coverage database in the `.d` files. It employs a single file, `tcov.lock`, for this purpose. Consequently, only one executable compiled with `-xa` (C) or `-a` (other compilers) should be running on the system at a time. If the execution of the program compiled with the `-xa` (or `-a`) option is manually terminated, then the `tcov.lock` file must be deleted manually.

Files compiled with the `-xa` or `-a` option call the profiling tool functions automatically when a program is linked for `tcov` profiling. At program exit, these functions combine the information collected at runtime for file `xyz.f` (for example) with the existing profiling information stored in file `xyz.d`. To ensure this information is not corrupted by several people simultaneously running a profiled binary, a `xyz.d.lock` lock file is created for `xyz.d` for the duration of the update. If there are any errors in opening or reading `xyz.d` or its lock file, or if there are inconsistencies between the runtime information and the stored information, the information stored in `xyz.d` is not changed.

If you edit and recompile `xyz.f` the number of counters in `xyz.d` can change. This is detected if an old profiled binary is run.

If too many people are running a profiled binary, some of them cannot obtain a lock. An error message is displayed after a delay of several seconds. The stored information is not updated. This locking is safe across a network. Since locking is performed on a file-by-file basis, other files may be correctly updated.

The profiling functions attempt to deal with automounted file systems that have become inaccessible. They still fail if the file system containing a coverage data file is mounted with different names on different machines, or if the user running the profiled binary does not have permission to write to either the coverage data file or the directory containing it. Be sure all the directories are uniformly named and writable by anyone expected to run the binary.

## Errors Reported by `tcov` Runtime Functions

The following error messages may be reported by the `tcov` runtime functions:

- The user running the binary lacks permission to read or write to the coverage data file. The problem also occurs if the coverage data file has been deleted.

```
tcov_exit: Could not open coverage data file 'coverage-data-file-name'  
because 'system-error-message-string' .
```

- The user running the binary lacks permission to write to the directory containing the coverage data file. The problem also occurs if the directory containing the coverage data file is not mounted on the machine where the binary is being run.

```
tcov_exit: Could not write coverage data file 'coverage-data-file-name'  
because 'system-error-message-string' .
```

- Too many users are trying to update a coverage data file at the same time. The problem also occurs if a machine has crashed while a coverage data file is being updated, leaving behind a lock file. In the event of a crash, the longer of the two files should be used as the post-crash coverage data file. Manually remove the lock file.

```
tcov_exit: Failed to create lock file 'lock-file-name' for coverage  
data file 'coverage-data-file-name' after 5 tries. Is someone else  
running this executable?
```

- No memory is available, and the standard I/O package will not work. You cannot update the coverage data file at this point.

```
tcov_exit: Stdio failure, probably no memory left.
```

- The lock file name is longer by six characters than the coverage data file name. Therefore, the derived lock file name may not be legal.

```
tcov_exit: Coverage data file path name too long (length  
characters) 'coverage-data-file-name' .
```

- A library or binary that has `tcov` profiling enabled is simultaneously being run, edited, and recompiled. The old binary expects a coverage data file of a certain size, but the editing often changes that size. If the compiler creates a new

coverage data file at the same time that the old binary is trying to update the old coverage data file, the binary may see an apparently empty or corrupt coverage file.

```
tcov_exit: Coverage data file 'coverage-data-file-name' is too short.  
Is it out of date?
```

---

## Using `tcov` Enhanced for Statement-Level Analysis

Like the original `tcov`, `tcov` Enhanced gives line-by-line information on how a program executes. It produces a copy of the source file, annotated to show which lines are used and how often. It also gives a summary of information about basic blocks. `tcov` Enhanced works with both C and C++ source files.

`tcov` Enhanced overcomes some of the shortcomings of the original `tcov`. The improved features of `tcov` Enhanced are:

- It provides more complete support for C++.
- It supports code found in `#include` header files and corrects a flaw that obscured coverage numbers for template classes and functions.
- Its runtime is more efficient than the original `tcov` runtime.
- It is supported for all the platforms that the compilers support.

To generate annotated source code using `tcov` Enhanced:

- 1. Compile your program with the `-xprofile=tcov` compiler option.**

Unlike `tcov`, `tcov` Enhanced does not generate any files at compile time.

- 2. Run your program.**

A directory is created to store the profile data, and a single coverage data file called `tcovd` is created in that directory. By default, the directory is created in the location where you run the program *program-name*, and it is called *program-name.profile*. The directory is also known as the *profile bucket*. The defaults can be changed using environment variables (see [“`tcov` Directories and Environment Variables” on page 169](#)).

### 3. Run `tcov` to generate annotated source code.

The syntax of the `tcov` command is as follows.

```
% tcov option-list source-file-list
```

Here, *source-file-list* is a list of the source code filenames, and *option-list* is a list of options, which can be obtained from the `tcov(1)` man page. You must include the `-x` option to enable `tcov` Enhanced processing.

The default output of `tcov` Enhanced is a set of annotated source files whose names are derived by appending `.tcov` to the corresponding source file name.

The following example illustrates the syntax of `tcov` Enhanced.

```
% cc -xprofile=tcov -o index.assist index.assist.c
% index.assist
% tcov -x index.assist.profile index.assist.c
```

The output of `tcov` Enhanced is identical to the output from the original `tcov`.

## Creating Profiled Shared Libraries for `tcov` Enhanced

You can create profiled shared libraries for use with `tcov` Enhanced by including the `-xprofile=tcov` compiler option, as shown in the following example.

```
% cc -G -xprofile=tcov -o foo.so.1 foo.o
```

## Locking Files

`tcov` Enhanced uses a simple file-locking mechanism for updating the block coverage data file. It employs a single file created in the same directory as the `tcovd` file. The file name is `tcovd.temp.lock`. If execution of the program compiled for coverage analysis is manually terminated, then the lock file must be deleted manually.



The locking scheme does an exponential back-off if there is a contention for the lock. If, after five tries, the `tcov` runtime cannot acquire the lock, it exits, and the data is lost for that run. In this case, the following message is displayed.

```
tcov_exit: temp file exists, is someone else running this
executable?
```

## `tcov` Directories and Environment Variables

When you compile a program for `tcov` and run the program, the running program generates a profile bucket. If a previous profile bucket exists, the program uses that profile bucket. If a profile bucket does not exist, it creates the profile bucket.

The profile bucket specifies the directory where the profile output is generated. The name and location of the profile output are controlled by defaults that you can modify with environment variables.

---

**Note** – `tcov` uses the same defaults and environment variables that are used by the compiler options that you use to gather profile feedback: `-xprofile=collect` and `-xprofile=use`. For more information about these compiler options, see the documentation for the relevant compiler.

---

The default profile bucket is named after the executable with a `.profile` extension and is created in the directory where the executable is run. Therefore, if you run a program called `/usr/bin/xyz` from `/home/userdir`, the default behavior is to create a profile bucket called `xyz.profile` in `/home/userdir`.

A UNIX process can change its current working directory during the execution of a program. The current working directory used to generate the profile bucket is the current working directory of the program at exit. In the rare case where a program actually does change its current working directory during execution, you can use the environment variables to control where the profile bucket is generated.

You can set the following environment variables to modify the defaults:

- `SUN_PROFDATA`

Can be used to specify the name of the profile bucket at runtime. The value of this variable is always appended to the value of `SUN_PROFDATA_DIR` if both variables are set. Doing this may be useful if the name of the executable is not the same as the value in `argv[0]` (for example, the invocation of the executable was through a symbolic link with a different name).

- SUN\_PROFDATA\_DIR

Can be used to specify the name of the directory that contains the profile bucket. It is used at runtime and by the `tcov` command.

- TCOVDIR

TCOVDIR is supported as a synonym for SUN\_PROFDATA\_DIR to maintain backward compatibility. Any setting of SUN\_PROFDATA\_DIR causes TCOVDIR to be ignored. If both SUN\_PROFDATA\_DIR and TCOVDIR are set, a warning is displayed when the profile bucket is generated.

TCOVDIR is used at runtime and by the `tcov` command.

# Index

---

## A

- accessible documentation, 19
- `addpath` command, 95
- address spaces, text and data regions, 130
- aliased functions, 131
- alternate entry points in Fortran functions, 132
- Analyzer, *See* Performance Analyzer
- annotated disassembly code, *See* disassembly code, annotated
- annotated source code, *See* source code, annotated
- API, Collector, 48
- appending path to files, 95
- `arc`, call graph, defined, 158
- archiving load objects in experiments, 69, 75
- asynchronous I/O library, interaction with data collection, 48
- attaching the Collector to a running process, 76
- attributed metrics
  - defined, 38
  - effect of recursion on, 41
  - illustrated, 40
  - use of, 39

## B

- body functions, compiler-generated
  - defined, 124
  - displayed by the Performance Analyzer, 134
  - names, 125
  - propagation of inclusive metrics, 128

## C

- C++ name demangling, setting default library in `.er.rc` file, 103
- call stacks
  - default alignment and depth in the Timeline tab, 104
  - defined, 117
  - effect of tail-call optimization on, 120
  - incomplete unwind, 129
  - mapping addresses to program structure, 130
  - unwinding, 117
- callers-callees metrics
  - attributed, defined, 38
  - displaying list of in `er_print`, 100
  - printing for a single function in `er_print`, 91
  - printing in `er_print`, 90
  - selecting in `er_print`, 90
  - sort order in `er_print`, 91
- clock-based profiling
  - accuracy of metrics, 114
  - collecting data in `dbx`, 71
  - collecting data with `collect`, 63
  - comparison with `gethrtime` and `gethrvtime`, 113
  - data in profile packet, 111
  - defined, 29
  - distortion due to overheads, 113
  - interval, *See* profiling interval
  - metrics, 29, 112
- cloned functions, 133
- `collect` command
  - archiving (`-A`) option, 69
  - clock-based profiling (`-p`) option, 63

- collecting data with, 62
- data collection options, 62
- data limit (-L) option, 69
- dry run (-n) option, 69
- experiment control options, 65
- experiment directory (-d) option, 68
- experiment group (-g) option, 68
- experiment name (-o) option, 68
- follow descendant processes (-F) option, 65
- hardware-counter overflow profiling (-h) option, 63
- heap tracing (-H) option, 64
- Java version (-j) option, 66
- listing the options of, 62
- miscellaneous options, 69
- MPI tracing (-m) option, 65
- output options, 68
- pause and resume data recording (-y) option, 67
- periodic sampling (-S) option, 65
- readme display (-R) option, 69
- record sample point (-l) option, 67
- stop target after exec (-x) option, 67
- synchronization wait tracing (-s) option, 64
- syntax, 62
- verbose (-v) option, 70
- version (-V) option, 70
- with `ppgsz` command, 82

## Collector

- API, using in your program, 48, 50
- attaching to a running process, 76
- defined, 24, 28
- disabling in `dbx`, 74
- enabling in `dbx`, 74
- running in `dbx`, 70
- running with `collect`, 62

`collectorAPI.h`, 50

common subexpression elimination, 142

compiler commentary

- classes defined, 93
- description of, 141
- selecting for annotated disassembly listing in `er_print`, 94
- selecting for annotated source listing in `er_print`, 93

compiler-generated body functions

- defined, 124
- displayed by the Performance Analyzer, 134
- names, 125

- propagation of inclusive metrics, 128

compilers, accessing, 15

compiling

- affects of optimization on program analysis, 44
- debug symbol information format, 44
- effect of static linking on data collection, 44
- for `gprof`, 158
- for Lines analyses, 44
- for `prof`, 156
- for `tcov`, 161
- for `tcov Enhanced`, 167
- Java programming language, 45
- linking for data collection, 44
- source code for annotated Source and Disassembly, 44
- static linking of libraries, 44

copying an experiment, 150

correlation, effect on metrics, 113

CPUs

- listing selected, in `er_print`, 97
- selecting in `er_print`, 98

## D

data collection

- controlling from your program, 48
- disabling from your program, 52
- disabling in `dbx`, 74
- dynamic memory allocation effects, 45
- enabling in `dbx`, 74
- from MPI programs, 78
- linking for, 44
- MPI program, using `collect`, 81
- MPI program, using `dbx`, 81
- pausing for `collect`, 67
- pausing from your program, 51
- pausing in `dbx`, 74
- preparing your program for, 45
- program control of, 48
- rate of, 60
- resuming for `collect`, 67
- resuming from your program, 51
- resuming in `dbx`, 74
- segmentation faults, 46
- using `collect`, 62
- using `dbx`, 70

data types, 28

- clock-based profiling, 29
- default, in the Timeline tab, 104
- hardware counter overflow profiling, 30
- heap tracing, 35
- MPI tracing, 35
- synchronization wait tracing, 33
- data\_objects command, 95
- data\_olayout command, 96
- data\_osingle command, 95
- dataobject
  - <Scalar> descriptor, 139
  - <Total> descriptor, 139
  - defined, 137
  - in hardware counter experiment, 95
  - scope, 138
- dbx
  - collecting data under MPI, 81
  - running the Collector in, 70
- dbx collector subcommands
  - archive, 75
  - dbxsample, 73
  - disable, 74
  - enable, 74
  - enable\_once (obsolete), 76
  - hwprofile, 71
  - limit, 75
  - pause, 74
  - profile, 71
  - quit (obsolete), 76
  - resume, 74
  - sample, 73
  - sample record, 74
  - show, 76
  - status, 76
  - store, 75
  - store filename (obsolete), 76
  - synctrace, 72, 73
- defaults
  - setting in a defaults file, 102
- descendant processes
  - collecting data for all followed, 65
  - collecting data for selected, 76
  - experiment location, 58
  - experiment names, 59
  - followed by Collector, 56
  - limitations on data collection for, 56
- directives, parallelization
  - attribution of metrics to, 142

- microtasking library calls from, 124
- disassembly code, annotated
  - description, 142
  - for cloned functions, 133
  - hardware counter metric attribution, 147
  - instruction issue dependencies, 144
  - interpreting, 144
  - location of executable, 60
  - metric formats, 141
  - printing in er\_print, 93
  - setting preferences in er\_print, 94
  - setting the highlighting threshold in er\_print, 94
  - viewing with er\_src, 150
- disk space, estimating for experiments, 60
- documentation index, 18
- documentation, accessing, 18 to 19
- dynamically compiled functions
  - Collector API for, 52
  - definition, 135

## E

- entry points, alternate, in Fortran functions, 132
- environment variables
  - JAVA\_PATH, 56
  - JDK\_1\_4\_HOME, 56
  - JDK\_HOME, 56
  - LD\_LIBRARY\_PATH, 78
  - LD\_PRELOAD, 78
  - PATH, 56
  - SUN\_PROFDATA, 169
  - SUN\_PROFDATA\_DIR, 170
  - TCOVDIR, 161, 170
- er\_archive utility, 152
- er\_cp utility, 150
- er\_export utility, 153
- er\_mv utility, 150
- er\_print commands
  - addpath, 95
  - allocs, 91
  - callers-callees, 90
  - cmetric\_list, 100
  - cmetrics, 90
  - cpu\_list, 97
  - cpu\_select, 98

- csingle, 91
- csort, 91
- data\_objects, 95
- data\_olayout, 96
- data\_osingle, 95
- dcc, 94
- disasm, 93
- dmetrics, 102
- dsort, 103
- exp\_list, 96
- fsingle, 89
- fsummary, 89
- functions, 87
- gdemangle, 103
- header, 101
- help, 105
- javamode, 101
- leaks, 91
- limit, 100
- lines, 92
- lsummary, 92
- lwp\_list, 97
- lwp\_select, 98
- mapfile, 105
- metric\_list, 100
- metrics, 88
- name, 100
- object\_list, 99
- object\_select, 99
- objects, 101
- outfile, 100
- overview, 101
- pcs, 92
- psummary, 92
- quit, 105
- sample\_list, 96
- sample\_select, 98
- scc, 93
- script, 105
- setpath, 95
- sort, 88
- source, 92
- src, 92
- statistics, 101
- sthresh, 94
- thread\_list, 97
- thread\_select, 98
- tlldata, 104
- tlmode, 104
- Version, 105
- version, 105
- er\_print utility
  - command-line options, 84
  - commands, *See er\_print* commands
  - metric keywords, 86
  - metric lists, 84
  - purpose, 83
  - syntax, 84
- er\_rm utility, 150
- er\_src utility, 150
- errors reported by tcov, 165
- events
  - default display type in the Timeline tab, 104
- exclusive metrics
  - defined, 38
  - for PLT instructions, 118
  - how computed, 117
  - illustrated, 40
  - use of, 39
- execution statistics
  - comparison of times with the <Total> function, 114
  - printing in er\_print, 101
- experiment directory
  - default, 58
  - specifying in dbx, 75
  - specifying with collect, 68
- experiment groups
  - default name, 58
  - defined, 58
  - name restrictions, 59
  - removing, 150
  - specifying name in dbx, 75
  - specifying name with collect, 68
- experiment names
  - default, 58
  - MPI default, 59, 80
  - MPI, using MPI\_comm\_rank and a script, 82
  - restrictions, 58
  - specifying in dbx, 75
  - specifying with collect, 68
- experiments
  - See also* experiment directory; experiment groups; experiment names
  - appending current path, 95
  - archiving load objects in, 69, 75

- copying, 150
- default name, 58
- defined, 58
- groups, 58
- header information in `er_print`, 101
- limiting the size of, 69, 75
- listing in `er_print`, 96
- location, 58
- moving, 59, 150
- moving MPI, 80
- MPI storage issues, 79
- naming, 58
- removing, 150
- setting mode for Java, 101
- setting path to find files, 95
- storage requirements, estimating, 60
- terminating from your program, 52
- where stored, 68, 75

explicit multithreading, 121

## F

- fast traps, 119
- Fortran
  - alternate entry points, 132
  - Collector API, 48
  - subroutines, 131
- frames, stack, *See* stack frames
- function calls
  - between shared objects, 118
  - imputed, in OpenMP programs, 127
  - in single-threaded programs, 118
  - recursive, metric assignment to, 41
- function list
  - printing in `er_print`, 87
  - sort order, specifying in `er_print`, 88
- function names, C++
  - choosing long or short form in `er_print`, 100
  - setting default demangling library in `.er.rc` file, 103
- function-list metrics
  - displaying list of in `er_print`, 100
  - selecting default in `.er.rc` file, 102
  - selecting in `er_print`, 88
  - setting default sort order in `.er.rc` file, 103
- functions
  - @plt, 118

- address within a load object, 131
- aliased, 131
- alternate entry points (Fortran), 132
- body, compiler-generated, *See* body functions, compiler-generated
- cloned, 133
- Collector API, 48, 52
- definition of, 131
- dynamically compiled, 52, 135
- global, 131
- inlined, 133
- MPI, traced, 35
- non-unique, names of, 131
- outline, 135
- static, in stripped shared libraries, 132
- static, with duplicate names, 131
- system library, interposition by Collector, 46
- <Total>, 136
- <Unknown>, 135
- variation in addresses of, 130
- wrapper, 132

## G

- gprof
  - limitations, 158
  - output from, interpreting, 159
  - summary, 155
  - using, 158

## H

- hardware counter library, `libcpc.so`, 55
- hardware counter list
  - description of fields, 31
  - obtaining with `collect`, 62
  - obtaining with `dbx` collector, 72
- hardware counters
  - choosing with `collect`, 63
  - choosing with `dbx` collector, 72
  - data objects and metrics, 95
  - list described, 31
  - obtaining a list of, 62, 72
  - overflow value, 30
- hardware-counter overflow profiling
  - collecting data with `collect`, 63

- collecting data with `dbx`, 71
- data in profile packet, 115
- defined, 30
- limitations, 55
- hardware-counter overflow value
  - consequences of too small or too large, 115
  - defined, 30
  - experiment size, effect on, 60
  - setting in `dbx`, 72
  - setting with `collect`, 64
- heap tracing
  - collecting data in `dbx`, 73
  - collecting data with `collect`, 64
  - metrics, 35
  - preloading the Collector library, 78
- high metric values
  - in annotated disassembly code, 94
  - in annotated source code, 94

## I

- inclusive metrics
  - defined, 38
  - effect of recursion on, 41
  - for PLT instructions, 118
  - how computed, 117
  - illustrated, 40
  - use of, 39
- inlined functions, 133
- input file
  - terminating in `er_print`, 105
  - to `er_print`, 105
- instruction issue
  - delay, 146
  - grouping, effect on annotated disassembly, 144
- intermediate files, use for annotated source listings, 129
- interposition by Collector on system library functions, 46
- interval, profiling, *See* profiling interval
- interval, sampling, *See* sampling interval

## J

- Java
  - dynamically compiled methods, 52, 135

- memory allocations, 34
- monitors, 34
- profiling limitations, 56
- setting `er_print` display output for, 101
- JAVA\_PATH environment variable, 56
- `javamode` command, 101
- JDK\_1\_4\_HOME environment variable, 56
- JDK\_HOME environment variable, 56
- JVM versions, 57

## K

- keywords, `metric`, `er_print` utility, 86

## L

- LD\_LIBRARY\_PATH environment variable, 78
- LD\_PRELOAD environment variable, 78
- leaf PC, defined, 117
- leaks, memory: definition, 35
- `libaio.so`, interaction with data collection, 48
- `libcollector.h`, 49
  - as part of C and C++ interface to collector, 49
  - as part of Java programming language interface to collector, 50
- `libcollector.so` shared library
  - preloading, 78
  - using in your program, 48
- `libcpc.so`, use of, 55
- `libfcollector.h`, 49
- libraries
  - `collectorAPI.h`, 50
  - interposition on, 46
  - `libaio.so`, 48
  - `libcollector.so`, 48, 78
  - `libcpc.so`, 47, 55
  - `libthread.so`, 46, 121, 125
  - MPI, 46, 78
  - static linking, 44
  - stripped shared, and static functions, 132
  - system, 46
- limitations
  - descendant process data collection, 56
  - experiment group names, 59
  - experiment name, 58



- hardware-counter overflow profiling, 55
  - Java profiling, 56
  - profiling interval value, 54
  - tcov, 161
  - limiting output in `er_print`, 100
  - limiting the experiment size, 69, 75
  - load objects
    - addresses of functions, 131
    - contents of, 130
    - defined, 130
    - listing selected, in `er_print`, 99
    - printing list in `er_print`, 101
    - selecting in `er_print`, 99
    - symbol tables, 130
    - writing layouts of, 96
  - lock file management
    - tcov, 165
    - tcov Enhanced, 168
  - LWPs
    - creation by threads library, 121
    - listing selected, in `er_print`, 97
    - selecting in `er_print`, 98
- M**
- man pages, accessing, 15
  - MANPATH environment variable, setting, 17
  - mapfiles
    - generating with `er_print`, 105
  - memory allocation effects on data collection, 45
  - memory allocations, 35
  - memory leaks, definition, 35
  - methods, *See* functions
  - metrics
    - attributed, *See* attributed metrics
    - clock-based profiling, 29, 112
    - defined, 27
    - effect of correlation, 113
    - exclusive, *See* exclusive metrics
    - function-list, *See* function-list metrics
    - hardware counter, attributing to
      - instructions, 147
    - heap tracing, 35
    - inclusive, *See* inclusive metrics
    - interpreting for instructions, 144
    - interpreting for source lines, 141
    - memory allocation, 35
    - MPI tracing, 36
    - synchronization wait tracing, 34
    - timing, 29
  - microstates
    - contribution to metrics, 112
    - switching, 119
  - microtasking library routines, 124
  - moving an experiment, 59, 150
  - MPI experiments
    - default name, 59
    - moving, 80
    - storage issues, 79
  - MPI programs
    - attaching to, 78
    - collecting data from, 78
    - collecting data with `collect`, 81
    - collecting data with `dbx`, 81
    - experiment names, 59, 79, 80
    - experiment storage issues, 79
  - MPI tracing
    - collecting data in `dbx`, 73
    - collecting data with `collect`, 65
    - data in profile packet, 116
    - functions traced, 35
    - interpretation of metrics, 116
    - metrics, 36
    - preloading the Collector library, 78
  - multithreaded applications
    - attaching the Collector to, 76
    - execution sequence, 125
  - multithreading
    - explicit, 121
    - parallelization directives, 124
- N**
- naming an experiment, 58
  - networked disks, 58
  - nfs, 58
  - non-unique function names, 131
- O**
- OpenMP parallelization, 124

- optimizations
  - common subexpression elimination, 142
  - program analysis affect of, 44
  - tail-call, 120
- options, command-line, `er_print` utility, 84
- outline functions, 135
- output file, in `er_print`, 100
- overflow value, hardware-counter, *See* hardware-counter overflow value
- overview data, printing in `er_print`, 101

## P

- parallel execution
  - call sequence, 125
  - directives, 124
- PATH environment variable, 56
- PATH environment variable, setting, 16
- path to files, 95
- pausing data collection
  - for `collect`, 67
  - from your program, 51
  - in `dbx`, 74
- PCs
  - defined, 117
  - from PLT, 118
  - ordered list in `er_print`, 92
- Performance Analyzer
  - defined, 24
- performance data, conversion into metrics, 27
- performance metrics, *See* metrics
- PLT (Program Linkage Table), 118, 148
  - @plt function, 118
- `ppgsz` command, 82
- preloading `libcollector.so`, 78
- printing the current path, 95
- process address-space text and data regions, 130
- `prof`
  - limitations, 158
  - output from, 157
  - summary, 155
  - using, 156
- profile bucket, `tcov Enhanced`, 167, 169
- profile packet
  - clock-based data, 111

- hardware-counter overflow data, 115
- MPI tracing data, 116
- size of, 60
- synchronization wait tracing data, 114
- profiled shared libraries, creating
  - for `tcov`, 164
  - for `tcov Enhanced`, 168
- profiling interval
  - defined, 29
  - experiment size, effect on, 60
  - limitations on value, 54
  - setting with `dbx collector`, 71
  - setting with the `collect` command, 63, 71
- profiling, defined, 28
- program counter (PC), defined, 117
- program execution
  - call stacks described, 117
  - explicit multithreading, 121
  - OpenMP parallel, 125
  - shared objects and function calls, 118
  - signal handling, 119
  - single-threaded, 118
  - tail-call optimization, 120
  - traps, 119
- Program Linkage Table (PLT), 118, 148
- program structure, mapping call-stack addresses to, 130

## R

- recursive function calls
  - apparent, in OpenMP programs, 128
  - metric assignment to, 41
- removing an experiment or experiment group, 150
- restrictions, *See* limitations
- resuming data collection
  - for `collect`, 67
  - from your program, 51
  - in `dbx`, 74

## S

- samples
  - circumstances of recording, 37
  - defined, 37
  - information contained in packet, 37

- interval, *See* sampling interval
- listing selected, in `er_print`, 96
- manual recording in `dbx`, 74
- manual recording with `collect`, 67
- periodic recording in `dbx`, 73
- periodic recording with `collect`, 65
- recording from your program, 51
- recording when `dbx` stops a process, 73
- selecting in `er_print`, 98
- Sampling Collector, *See* Collector
- sampling interval
  - defined, 37
  - setting in `dbx`, 73
  - setting with the `collect` command, 65
- <Scalar> dataobject descriptor, 139
- segmentation faults during data collection, 46
- setpath command, 95
- setuid, use of, 48
- shared objects, function calls between, 118
- shell prompts, 15
- signal handlers
  - installed by Collector, 47, 119
  - user program, 47
- signals
  - calls to handlers, 119
  - profiling, 47
  - profiling, passing from `dbx` to `collect`, 67
  - use for manual sampling with `collect`, 67
  - use for pause and resume with `collect`, 67
- single-threaded program execution, 118
- sort order
  - callers-callees metrics, in `er_print`, 91
  - function list, specifying in `er_print`, 88
- source code, annotated
  - compiler commentary, 141
  - description, 140
  - for cloned functions, 133
  - from `tcov`, 163
  - interpreting, 141
  - location of source files, 60
  - metric formats, 141
  - parallelization directives in, 142
  - printing in `er_print`, 92
  - setting compiler commentary classes in
    - `er_print`, 93
  - setting the highlighting threshold in
    - `er_print`, 94
  - <Unknown> line, 142
    - use of intermediate files, 129
    - viewing with `er_src`, 150
- source lines
  - ordered list in `er_print`, 92
- stack frames
  - defined, 118
  - from trap handler, 120
  - reuse of in tail-call optimization, 120
- static functions
  - duplicate names, 131
  - in stripped shared libraries, 132
- static linking, effect on data collection, 44
- storage requirements, estimating for
  - experiments, 60
- subroutines, *See* functions
- summary metrics
  - for a single function, printing in `er_print`, 89
  - for all functions, printing in `er_print`, 89
- SUN\_PROFDATA environment variable, 169
- SUN\_PROFDATA\_DIR environment variable, 170
- symbol tables, load-object, 130
- synchronization delay events
  - data in profile packet, 114
  - defined, 33
  - metric defined, 34
- synchronization wait time
  - defined, 33, 114
  - metric, defined, 34
  - with unbound threads, 115
- synchronization wait tracing
  - collecting data in `dbx`, 72
  - collecting data with `collect`, 64
  - data in profile packet, 114
  - defined, 33
  - metrics, 34
  - preloading the Collector library, 78
  - threshold, *See* threshold, synchronization wait
    - tracing
  - wait time, 33, 114
- syntax
  - `er_archive` utility, 152
  - `er_export` utility, 153
  - `er_print` utility, 84
  - `er_src` utility, 150

## T

tail-call optimization, 120

tcov

- annotated source code, 163
- compiling a program for, 161
- errors reported by, 165
- limitations, 161
- lock file management, 165
- output, interpreting, 163
- profiled shared libraries, creating, 164
- summary, 155
- using, 161

tcov Enhanced

- advantages of, 167
- compiling a program for, 167
- lock file management, 168
- profile bucket, 167, 169
- profiled shared libraries, creating, 168
- using, 167

TCOVDIR environment variable, 161, 170

threads

- bound and unbound, 121, 128
- creation of, 121
- library, 46, 121, 125
- listing selected, in `er_print`, 97
- main, 125
- scheduling of, 121, 124
- selecting in `er_print`, 98
- system, 114, 125
- wait mode, 128
- worker, 121, 125

threshold, highlighting

- in annotated disassembly code, `er_print`, 94
- in annotated source code, `er_print`, 94

threshold, synchronization wait tracing

- calibration, 33
- defined, 33
- effect on collection overhead, 115
- setting with `dbx` collector, 72
- setting with the `collect` command, 64, 73

TLB (translation lookaside buffer) misses, 120, 147

<Total> dataobject descriptor, 139

<Total> function

- comparing times with execution statistics, 114
- described, 136

traps, 119

typographic conventions, 14

## U

<Unknown> dataobject descriptor

- Unascertainable element, 139
- Unidentified element, 139
- Unresolvable element, 139
- Unspecific element, 139
- Unverifiable element, 139

<Unknown> function

- callers and callees, 136
- mapping of PC to, 135

<Unknown> line, in annotated source code, 142

unwinding the call stack, 117

## V

version information

- for `collect`, 70
- for `er_cp`, 150
- for `er_mv`, 150
- for `er_print`, 105
- for `er_rm`, 150
- for `er_src`, 151

## W

wait time, *See* synchronization wait time

wrapper functions, 132

## X

-xdebugformat

- setting debug symbol information format, 44