

Rational® PurifyPlus
Rational® Purify®
Rational® PureCoverage®
Rational® Quantify®

Installing and Getting Started

VERSION: 2003.06.00

PART NUMBER: 800-026184-000

UNIX

Legal Notices

©1992-2003, Rational Software Corporation. All rights reserved.

This manual (the "Work") is protected under the copyright laws of the United States and/or other jurisdictions, as well as various international treaties. Any reproduction or distribution of the Work is expressly prohibited without the prior written consent of Rational Software Corporation.

Part Number: 800-026184-000

Version Number: 2003.06.00

Rational, Rational Software Corporation, the Rational logo, Rational Developer Network, AnalystStudio, ClearCase, ClearCase Attache, ClearCase MultiSite, ClearDDTS, ClearGuide, ClearQuest, ClearTrack, Connexis, e-Development Accelerators, DDTS, Object Testing, Object-Oriented Recording, ObjecTime, ObjecTime Design Logo, Objectory, PerformanceStudio, PureCoverage, PureDDTS, PureLink, Purify, Quantify, Rational Apex, Rational CRC, Rational Process Workbench, Rational Rose, Rational Suite, Rational Suite ContentStudio, Rational Summit, Rational Visual Test, Rational Unified Process, RUP, RequisitePro, ScriptAssure, SiteCheck, SiteLoad, SoDA, TestFactory, TestFoundation, TestStudio, TestMate, VADS, and XDE, among others, are trademarks or registered trademarks of Rational Software Corporation in the United States and/or in other countries. All other names are used for identification purposes only, and are trademarks or registered trademarks of their respective companies.

Portions covered by U.S. Patent Nos. 5,193,180 and 5,335,344 and 5,535,329 and 5,574,898 and 5,649,200 and 5,675,802 and 5,754,760 and 5,835,701 and 6,049,666 and 6,126,329 and 6,167,534 and 6,206,584. Additional U.S. Patents and International Patents pending.

U.S. GOVERNMENT RIGHTS. All Rational software products provided to the U.S. Government are provided and licensed as commercial software, subject to the applicable license agreement. All such products provided to the U.S. Government pursuant to solicitations issued prior to December 1, 1995 are provided with "Restricted Rights" as provided for in FAR, 48 CFR 52.227-14 (JUNE 1987) or DFARS, 48 CFR 252.227-7013 (OCT 1988), as applicable.

WARRANTY DISCLAIMER. This document and its associated software may be used as stated in the underlying license agreement. Except as explicitly stated otherwise in such license agreement, and except to the extent prohibited or limited by law from jurisdiction to jurisdiction, Rational Software Corporation expressly disclaims all other warranties, express or implied, with respect to the media and software product and its documentation, including without limitation, the warranties of merchantability, non-infringement, title or fitness for a particular purpose or arising

from a course of dealing, usage or trade practice, and any warranty against interference with Licensee's quiet enjoyment of the product.

Third Party Notices, Code, Licenses, and Acknowledgements

Portions Copyright ©1992-1999, Summit Software Company. All rights reserved.

Microsoft, the Microsoft logo, Active Accessibility, Active Client, Active Desktop, Active Directory, ActiveMovie, Active Platform, ActiveStore, ActiveSync, ActiveX, Ask Maxwell, Authenticode, AutoSum, BackOffice, the BackOffice logo, bCentral, BizTalk, Bookshelf, ClearType, CodeView, DataTips, Developer Studio, Direct3D, DirectAnimation, DirectDraw, DirectInput, DirectX, DirectXJ, DoubleSpace, DriveSpace, FrontPage, Funstone, Genuine Microsoft Products logo, IntelliEye, the IntelliEye logo, IntelliMirror, IntelliSense, J/Direct, JScript, LineShare, Liquid Motion, Mapbase, MapManager, MapPoint, MapVision, Microsoft Agent logo, the Microsoft eMbedded Visual Tools logo, the Microsoft Internet Explorer logo, the Microsoft Office Compatible logo, Microsoft Press, the Microsoft Press logo, Microsoft QuickBasic, MS-DOS, MSDN, NetMeeting, NetShow, the Office logo, Outlook, PhotoDraw, PivotChart, PivotTable, PowerPoint, QuickAssembler, QuickShelf, RelayOne, Rushmore, SharePoint, SourceSafe, TipWizard, V-Chat, VideoFlash, Virtual Basic, the Virtual Basic logo, Visual C++, Visual C#, Visual FoxPro, Visual InterDev, Visual J++, Visual SourceSafe, Visual Studio, the Visual Studio logo, Vizact, WebBot, WebPIP, Win32, Win32s, Win64, Windows, the Windows CE logo, the Windows logo, Windows NT, the Windows Start logo, and XENIX, are either trademarks or registered trademarks of Microsoft Corporation in the United States and/or in other countries.

Sun, Sun Microsystems, the Sun Logo, Ultra, AnswerBook 2, medialib, OpenBoot, Solaris, Java, Java 3D, ShowMe TV, SunForum, SunVTS, SunFDDI, StarOffice, and SunPCi, among others, are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and other countries.

Purify is licensed under Sun Microsystems, Inc., U.S. Patent No. 5,404,499.

Licensee shall not incorporate any GLOBEtrouter software (FLEXIm libraries and utilities) into any product or application the primary purpose of which is software license management.

BasicScript is a registered trademark of Summit Software, Inc.

Design Patterns: Elements of Reusable Object-Oriented Software, by Erich Gamma, Richard Helm, Ralph Johnson and John Vlissides. Copyright © 1995 by Addison-Wesley Publishing Company, Inc. All rights reserved.

Additional legal notices are described in the legal_information.html file that is included in your Rational software installation.

Contents

Preface	ix
What's in this guide?	ix
Audience	x
Other resources	x
Contacting Rational technical publications	x
Contacting Rational technical support	x
1 Installing the products	1
Overview	1
Step 1: Obtaining a license for your Rational product	2
About licenses	2
Information you need to obtain a license	4
Obtaining a .upd import file using AccountLink	5
Step 2: Installing your Rational product	5
Information you need for rs_install	5
Installing the products using rs_install	15
Answers to questions about rs_install	16
Step 3: Post-installation configuration tasks	18
Checking and adjusting your configuration	21
Maintaining the rational.opt options file	22
Modifying the list of user names	23
Removing a previous product release	23
Requesting and installing the permanent or TLA license key	23
Requesting your permanent or TLA license key	24
Entering a permanent or TLA license key after initial installation	24
Supplemental notes	24
Creating an installation directory manually	24
Mounting the CD-ROM	25
Ejecting the CD-ROM	26
Using rs_install commands	27
Using the FLEXlm Software License Manager	27
2 Using Rational Purify	31
Rational Purify: What it does	31

Finding errors in Hello World.	32
Instrumenting a program.	33
Compiling and linking in separate stages	33
Running the instrumented program.	34
Seeing all your errors at a glance	35
Finding and correcting errors	36
Understanding the cause of the error	37
Correcting the ABR error	38
Finding leaked memory	39
Correcting the MLK error	40
Looking at the heap analysis	41
Comparing program runs	42
Suppressing Purify messages	42
Saving Purify output to a view file	43
Saving a run to a view file from the Viewer.	44
Opening a view file	44
Using your debugger with Purify	44
Using Purify with PureCoverage	45
Purify API functions.	45
Build-time options	46
Conversion characters for filenames.	47
Runtime options	47
Purify messages	49
How Purify finds memory-access errors	50
How Purify checks statically allocated memory	51
3 Using Rational PureCoverage	53
Rational PureCoverage: What it does.	53
Finding untested Java code	54
Finding untested C/C++ code	55
Instrumenting a C/C++ program	56
Running the instrumented C/C++ program.	56
Displaying C/C++ coverage data	57
Improving Hello World's test coverage	61
Viewing UNIX coverage data on Windows	64
Using report scripts	64
PureCoverage options.	66

Build-time options	66
Runtime options	67
Analysis-time options	67
Analysis-time mode options	68
4 Using Rational Quantify	69
Rational Quantify: What it does	69
How Quantify works: C/C++	69
How Quantify works: Java	70
Collecting performance data: C/C++	71
Interpreting the program summary: C/C++	72
Collecting performance data: Java	72
Interpreting the program summary: Java	73
Using Quantify's data analysis windows	74
The Function List window	75
Sorting the function list	75
Restricting functions	76
The Call Graph window	76
Using the pop-up menu	77
Expanding and collapsing descendants	78
The Function Detail window	78
Changing the scale and precision of data	79
Saving function detail data	79
The Annotated Source window	80
Changing annotations	80
Saving performance data on exit	81
Comparing program runs with qxdiff	81
Quantify options	82
Build-time options	82
qv runtime options	83
Runtime options	84
API functions: C/C++	85
API methods: Java	86
Index	89

Preface

What's in this guide?

This guide is designed to help you get up and running quickly with Rational® PurifyPlus, Purify®, PureCoverage®, and Quantify®. It includes information about:

- Installing the products.
Note: PurifyPlus is a Rational product that includes Purify, PureCoverage, and Quantify, and provides a unified procedure for installing all three applications on your system at the same time.
- Using Purify to pinpoint runtime errors and memory leaks everywhere in your C and C++ application code.
- Using PureCoverage to prevent untested C, C++, and Java applications from reaching end users. (PureCoverage can collect coverage data for Java applications running on a Solaris SPARC 32-bit Java virtual machine).
- Using Quantify to improve the performance of your C, C++, and Java applications by finding and eliminating bottlenecks. (Quantify can collect performance data for Java applications running on a Solaris SPARC 32-bit Java virtual machine).

Purify, PureCoverage, and Quantify—the essential tools for delivering high-performance UNIX applications—use patented Object Code Insertion (OCI) technology to instrument your program, inserting instructions into the program's object code. This enables you to check your entire program, including third-party code and shared libraries, even when you don't have the source code.

Starting to use Purify, PureCoverage, and Quantify is as easy as adding the product name (`purify`, `purecov`, or `quantify`) to the front of your link command line. For example:

```
% purify cc -g hello_world.c
```

Audience

Read this guide if you are responsible for installing Rational PurifyPlus, Purify, PureCoverage, or Quantify, or if you need an introduction to the use of Purify, PureCoverage, or Quantify.

Other resources

- A complete online help system is available for each application. Select **Help > Help topics**.

For help with a window, select **Help > On window**. For help with a specific menu item or control button in a window, select **Help > On context**, then click the menu item or control button.

Note: You can also view the help systems independently of the products. Open the following in your Netscape browser:

- `'purify -printhomedir'/UI/html/punix.htm`
- `'purecov -printhomedir'/UI/html/pcu.htm`
- `'quantify -printhomedir'/UI/html/qunix.htm`
- For information about Rational Software and Rational Software products, go to <http://www.rational.com>.

Contacting Rational technical publications

Please send any feedback about this documentation to the Rational technical publications department at techpubs@rational.com.

Contacting Rational technical support

You can contact Rational technical support by e-mail at support@rational.com.

You can also reach Rational technical support over the Web or by telephone. For contact information, as well as for answers to common questions about Purify, PureCoverage, and Quantify, go to <http://www.rational.com/support>.

Overview

The Rational® PurifyPlus for UNIX product family includes:

- Rational PurifyPlus
- Rational Purify®
- Rational PureCoverage®
- Rational Quantify®

PurifyPlus provides a license that allows you to use Purify, PureCoverage, and Quantify, as well as a unified procedure for installing all three applications on your system at the same time. Individual product licenses are also available.

Installing a product includes three steps:

- 1 Obtaining a license for the product.
- 2 Installing the product using the `rs_install` program.
- 3 Performing any necessary post-installation configuration procedures.

This chapter tells you how to gather the information you need to perform these steps, as well as basic instructions to get you started with each one. It also contains information about related tasks (such as uninstalling) and administering the GLOBEtrouter FLEXlm® Software License Manager that is included with your Rational Software product.

Note: *You may be able to complete licensing and installation without gathering any information at all.* If you are 1) using a startup or evaluation license, or a permanent license that has not expired, and 2) not setting up a new ClearQuest integration, just go ahead and run the `rs_install` program. The program will find existing licenses and perform a `Typical` installation (an installation of all the products for which you have a license, on the system where you're logged in) without requiring you to supply any information. Read *Installing the products using rs_install* on page 15.

Step 1: Obtaining a license for your Rational product

About licenses

When you purchase PurifyPlus, Purify, PureCoverage, or Quantify, you purchase a specific number of licenses for each product. The license key can be either a permanent license key, which has no expiration date, or a term license agreement (TLA) license key. Rational Software issues you a license key for the product that corresponds to the number and type of licenses you purchased. You need this license key to use the software.

The products also come with a startup license that you can use to get started using the product if, at the time of installation, you don't have the key you purchased. You can then request your key from AccountLink, Rational Software's license management web tool, and install it to ensure continued use of the product. The startup license key and other licensing information is available from the License Key Certificate included in the product packaging.

Evaluation licenses for the products are also available. These permit you to try the products for a limited period, ordinarily two weeks, so that you can see the benefits of using them.

Purify, PureCoverage, and Quantify use the FLEXlm Software License Manager from GLOBETrotter Software, Inc. to manage product licenses. For more information on FLEXlm, see *Using the FLEXlm Software License Manager* on page 27.

The following table provides definitions of license types, and helps you determine whether you need to perform Step 1, based on the type of license you are setting up.

License Type	Description	Instructions
Permanent	Allows the use of the product without time limits.	If you are upgrading from an earlier version of the product and still have a valid permanent or TLA license, go right to <i>Step 2: Installing your Rational product</i> on page 5.
Term License Agreement (TLA)	Allows the use of the product for a specific period of time.	<p>If you do not yet have a valid permanent or TLA license set up, gather the information specified in <i>Information you need to obtain a license</i> on page 4 and then request a license from AccountLink, Rational Software’s license management web tool. This applies if you are a first-time purchaser of the Rational PurifyPlus products.</p> <p>To enter a permanent or TLA license key after you’ve been using a startup or evaluation license, see <i>Entering a permanent or TLA license key after initial installation</i> on page 24.</p>
Startup	Is provided to get you up and running as soon as you receive your Rational product, if your permanent or TLA license key is not available. Note that you will have to enter a permanent or TLA key later to continue using the product, ordinarily within 90 days.	<p>Go right to <i>Step 2: Installing your Rational product</i> on page 5.</p> <p>When you receive your permanent or TLA license, see <i>Entering a permanent or TLA license key after initial installation</i> on page 24.</p>
Evaluation	Allows you to try out the product before purchasing for a limited time, ordinarily two weeks. Note that you will have to enter a permanent or TLA key later to continue using the product.	

Information you need to obtain a license

In addition to contact information, you will need to provide the following to AccountLink, Rational Software's license management web tool, in order to obtain a .upd import file. The .upd file contains licensing information and you will use it later in the licensing procedure.

Data for AccountLink	Notes	Your Entry
Your Rational account number	Source: your Rational license key certificate.	
License Type	Source: your Rational license key certificate Permanent and TLA licenses for PurifyPlus can be either Floating licenses or Named User licenses. Permanent and TLA licenses for Purify, PureCoverage, and Quantify are always Named User licenses.	
Rational Product Line	Source: your Rational license key certificate	DEVELOPER TOOLS
Product Name	Source: your Rational license key certificate.	PurifyPlus for UNIX, Purify for UNIX, PureCoverage for UNIX, or Quantify for UNIX
Quantity	Source: your Rational license key certificate. This is the number of Floating licenses or Named User licenses that you intend to install. This cannot be greater than the number of licenses you purchased.	
Host Name and hostid	This is the name and hostid of the machine that you intend to use as your license server host. If the license server host is different from the installation machine, you must have remote shell access from the installation machine to the license server host during installation. In addition, the installation directory must be accessible from the license server host. If you do not know the hostid of your license server host, you can download and run the tool <code>get_hostinfo.sh</code> after you have begun your license request on AccountLink; go to <i>Obtaining a .upd import file using AccountLink</i> on page 5.	

Obtaining a .upd import file using AccountLink

Access AccountLink at <http://www.rational.com/accountlink>. The AccountLink website provides instructions for requesting licenses.

When you supply the required information to AccountLink, you will receive by return email a file named

`license_for_<server name>.upd`. Save this file as a text file in a location that is accessible from the installation machine.

Go on to *Step 2: Installing your Rational product* on page 5.

Step 2: Installing your Rational product

Install your product using the script `rs_install`.

Information you need for `rs_install`

The information you need to supply to the `rs_install` script depends on what type of license you have, and on how you are setting it up.

- If you already have a valid license or are using a temporary or evaluation license, and if you are not setting up a new ClearQuest integration, and if you intend to perform a `Typical` installation (with the product or products installed on the same system as the system where you are logged in), you do not need to gather any additional information. Go right to *Installing the products using `rs_install`* on page 15.
- If you are installing a new permanent or TLA license, and you have a `.upd` file that you can use to import licensing data, go to page 6.
- If you are installing a new permanent or TLA license, and you **do not** have a `.upd` file that you can use to import licensing data, go to page 9.
- If you already have a permanent or TLA license set up, but are not performing a `Typical` installation, go to page 11.
- If you are installing a temporary or evaluation license, but are not performing a `Typical` installation, go page 13.

If you are installing a new permanent or TLA license and are importing a .upd file

Gather the information specified in the following table:

Data for rs_install	Notes	Your Entry
<p>The full pathname to the installation location (referred to in this chapter as Rational).</p>	<p>This is the directory for installing all Rational Software products. You must have 20 megabytes of free disk space for each installation of Purify, Quantify, and PureCoverage. You must have 60 megabytes for a complete installation of PurifyPlus. Note that these are per-platform figures.</p> <p>The directory must be accessible from every machine on which you plan to run the Rational products—both the machines on which users <i>instrument</i> their applications, and the machines on which users <i>run</i> their applications. The directory must be the same for each machine, so you cannot use a local automount path like <code>/tmp_mnt/Rational</code>.</p> <p>If <code>Rational</code> does not already exist, the installation program will create it when you enter the full pathname.</p> <p>If you are installing on a read-only file system, or want to create this directory manually, see <i>Supplemental notes</i> on page 24. This section also shows the structure of the directory after installation.</p>	

Data for rs_install	Notes	Your Entry
<p>Note: If you are installing a Named User license, you must supply user names for each individual who will be using the product. You must include your user name in order to perform the post-installation self-test successfully. User names are recorded in the FLEXlm options file, <code>rational.opt</code>. For information about the options file, see <i>Maintaining the rational.opt options file</i> on page 22.</p> <p>You do not need user names if you are installing a Floating license for PurifyPlus.</p> <p>To input user names, you need the data in A or B below, or the data in A supplemented with the data in B.</p>		
<p>A. Path of the PureLA directory containing the file <code>users.purela</code> (available only if you licensed an earlier version of the product using PureLA License Advisor).</p>	<p>If you are currently running the product under a PureLA license, you have the option of importing the user names from the PureLA database instead of entering them manually. The PureLA directory is located in the same parent directory as the previous product installation, which you can find with the command <code><product> -printhomedir</code>, where <code><product></code> is <code>purify</code>, <code>purecov</code>, or <code>quantify</code>.</p> <p>You can modify the list of imported user names, either while you're running <code>rs_install</code> or afterwards. If the number of user names is not the same as the number of licenses you bought, <code>rs_install</code> will help you correct the list.</p>	
<p>B. user names (all names; or some or none, in combination with an option to generate dummy names).</p>	<p>You can enter all user names. The number of names you enter must match the number of licenses you purchased.</p> <p>You can enter some user names, and then enter <code>-n</code> to populate the rest of the options file with dummy names as placeholders that you can replace later.</p> <p>Or you can just enter <code>-n</code> to enter nothing but dummy names, and update the options file later.</p>	
<p>Name for the Rational license file (<code><server name>.dat</code>), including full pathname.</p>	<p>The <code>rs_install</code> program, which creates this file when it runs, will suggest <code><server name>.dat</code> as a default.</p> <p>See <i>The Rational license file</i> on page 28 for information about the file.</p>	
<p>Note: If you have installed Rational ClearQuest, Rational Software's change request management system, it is possible for Purify and PureCoverage users to file change requests from within the Purify and PureCoverage graphical interface. To configure this integration between the products:</p> <ul style="list-style-type: none"> ▪ You as the installer can provide <code>rs_install</code> with site-wide default values. Note that individual users can override the default values if they have different requirements, following the instructions in the Purify and PureCoverage online help systems. (See "ClearQuest" in the online help index.) ▪ Alternatively, you can choose not to enter default values. In this case, users must enter their own values in order to use the integration. <p>The source for the following configuration values is your ClearQuest system administrator.</p>		

Data for rs_install	Notes	Your Entry
ClearQuest client interface.	<p>If you are installing on a Solaris or HP-UX platform, you can specify the default ClearQuest client for your users:</p> <ul style="list-style-type: none"> ▪ x for the ClearQuest native X client ▪ web for the ClearQuest web client <p>All other platforms support only the web client.</p> <p>Users must have ClearQuest on their path in order to use the native X client.</p>	
Name of the default ClearQuest database.	<p>Enter the name of a site-wide default ClearQuest database for your users who access ClearQuest through the native X client.</p> <p>Even if you have chosen web as your default interface, you may want to enter a value here for users who prefer to use the ClearQuest native X client.</p> <p>Source: ClearQuest system administrator.</p>	
URL for the ClearQuest web client.	<p>Specify the URL if you selected web as your default ClearQuest client interface.</p> <p>Even if you have chosen x as your default interface, you may want to enter a value here for users who prefer to use the web interface.</p> <p>Source: ClearQuest system administrator.</p>	

When you have gathered the information you need, go on to *Installing the products using rs_install* on page 15.

If you are installing a new permanent or TLA license without importing a .upd file

Gather the information specified in the following table:

Data for <code>rs_install</code>	Notes	Your Entry
<p>The full pathname to the installation location (referred to in this chapter as <code>Rational</code>).</p>	<p>This is the directory for installing all Rational Software products. You must have 20 megabytes of free disk space for each installation of Purify, Quantify, and PureCoverage. You must have 60 megabytes for a complete installation of PurifyPlus. Note that these are per-platform figures.</p> <p>The directory must be accessible from every machine on which you plan to run the Rational products—both the machines on which users <i>instrument</i> their applications, and the machines on which users <i>run</i> their applications. It must be the same for each machine, so you cannot use a local automount path like <code>/tmp_mnt/Rational</code>.</p> <p>If <code>Rational</code> does not already exist, the installation program will create it when you enter the full pathname.</p> <p>If you are installing on a read-only file system, or if you want to create this directory manually, see <i>Supplemental notes</i> on page 24. This section also shows you the structure of the directory after installation.</p>	
<p>Host name or IP address of the host machine on which the license server is to run (the "license server host").</p>	<p>If the license server host is different from the installation machine, you should have remote shell access from the installation machine to the license server host. If you do you have remote shell access, the <code>rs_install</code> program provides instructions for how to proceed.</p> <p>In addition, the installation directory must be accessible from the license server host.</p>	
<p>License server port number.</p>	<p>This is the port at which the license server listens for license requests. Default is 27000. You can use any port number that is not already in use. The <code>/etc/services</code> file on the license host lists all ports in use by most commonly used services, but other ports may be in use on your system as well. FLEXlm reserves ports 27000–27009 for its use; these ports are ordinarily available unless a different FLEXlm server on the license host is using them.</p> <p>The <code>rs_install</code> program checks to make sure that the license server port number does not conflict with entries in the <code>/etc/services</code> file on the license server host, or with NIS services.</p>	
<p>License quantity.</p>	<p>Source: your Rational license key certificate.</p>	

Data for rs_install	Notes	Your Entry
<p>Note: If you are installing a Named User license, you must supply user names for each individual who will be using the product. You must include your user name in order to perform the post-installation self-test successfully. User names are recorded in the FLEXlm options file, <code>rational.opt</code>. For information about the options file, see <i>Maintaining the rational.opt options file</i> on page 22.</p> <p>You do not need user names if you are installing a Floating license for PurifyPlus.</p> <p>To input user names, you need the data in A or B below, or the data in A supplemented with the data in B.</p>		
<p>A. Path of the PureLA directory containing the file <code>users.purela</code> (available only if you licensed an earlier version of the product using PureLA License Advisor).</p>	<p>If you are currently running the product under a PureLA license, you have the option of importing the user names from the PureLA database instead of entering them manually. The PureLA directory is located in the same parent directory as the previous product installation, which you can find with the command <code><product> -printhomedir</code>, where <code><product></code> is <code>purify</code>, <code>purecov</code>, or <code>quantify</code>.</p> <p>You can modify the list of imported user names, either while you're running <code>rs_install</code> or afterwards. If the number of user names is not the same as the number of licenses you bought, <code>rs_install</code> will help you correct the list.</p>	
<p>B. user names (all names; or some or none, in combination with an option to generate dummy names).</p>	<p>You can enter all user names. The number of names you enter must match the number of licenses you purchased.</p> <p>You can enter some user names, and then enter <code>-n</code> to populate the rest of the options file with dummy names as placeholders that you can replace later.</p> <p>Or you can just enter <code>-n</code> to enter nothing but dummy names, and update the options file later.</p>	
<p>Name for the Rational license file (<code><server name>.dat</code>), including full pathname.</p>	<p>The <code>rs_install</code> program, which creates this file when it runs, will suggest <code><server name>.dat</code> as a default.</p> <p>If you want to use an existing Rational license <code>.dat</code> file, enter its name, including full pathname, instead of the default. The <code>rs_install</code> program makes a backup of the existing license file before it processes the file with the new data. For information, see <i>The Rational license file</i> on page 28.</p>	
<p>Note: If you have installed Rational ClearQuest, Rational Software's change request management system, it is possible for Purify and PureCoverage users to file change requests from within the Purify and PureCoverage graphical interface. To configure this integration between the products:</p> <ul style="list-style-type: none"> ▪ You as the installer can provide <code>rs_install</code> with site-wide default values. Note that individual users can override the default values if they have different requirements, following the instructions in the Purify and PureCoverage online help systems. (See "ClearQuest" in the online help index.) ▪ Alternatively, you can choose not to enter default values. In this case, users must enter their own values in order to use the integration. <p>The source for the following configuration values is your ClearQuest system administrator.</p>		

Data for rs_install	Notes	Your Entry
ClearQuest client interface.	<p>If you are installing on a Solaris or HP-UX platform, you can specify the default ClearQuest client for your users:</p> <ul style="list-style-type: none"> ▪ x for the ClearQuest native X client ▪ web for the ClearQuest web client <p>All other platforms support only the web client.</p> <p>Users must have ClearQuest on their path in order to use the native X client.</p>	
Name of the default ClearQuest database.	<p>Enter the name of a site-wide default ClearQuest database for your users who access ClearQuest through the native X client.</p> <p>Even if you have chosen web as your default interface, you may want to enter a value here for users who prefer to use the ClearQuest native X client.</p> <p>Source: ClearQuest system administrator.</p>	
URL for the ClearQuest web client.	<p>Specify the URL if you selected web as your default ClearQuest client interface.</p> <p>Even if you have chosen x as your default interface, you may want to enter a value here for users who prefer to use the web interface.</p> <p>Source: ClearQuest system administrator.</p>	

When you have gathered the information you need, go on to *Installing the products using rs_install* on page 15.

If you already have a permanent or TLA license set up

If you are upgrading from an earlier version of the product and still have a valid permanent license, most of the information you need is already available in your system.

Gather the information specified in the following table:

Data for rs_install	Notes	Your Entry
Full pathname of your license file, <server name>.dat.	Alternatively, you can supply the host name or IP address of the host machine on which the license server is running (the “license server host”) and the license server port number.	

Data for <code>rs_install</code>	Notes	Your Entry
<p>The full pathname to the current installation location, or to the new installation location (referred to in this chapter as <code>Rational</code>).</p>	<p>This is the directory for installing all Rational Software products. You can use your existing installation; but if you use a new directory, it must meet the following requirements:</p> <ul style="list-style-type: none"> ▪ You must have 20 megabytes of free disk space for each installation of Purify, Quantify, and PureCoverage. You must have 60 megabytes for a complete installation of PurifyPlus. Note that these are per-platform figures. ▪ The directory must be accessible from every machine on which you plan to run the Rational products—both the machines on which users <i>instrument</i> their applications, and the machines on which users <i>run</i> their applications. It must be the same for each machine, so you cannot use a local automount path like <code>/tmp_mnt/Rational</code>. <p>If the directory you specify does not already exist, the installation program will create it when you enter the full pathname.</p> <p>If you are installing on a read-only file system, or if you want to create this directory manually, see <i>Supplemental notes</i> on page 24. This section also shows you the structure of the directory after installation.</p>	
<p>Note: If you have installed Rational ClearQuest, Rational Software's change request management system, it is possible for Purify and PureCoverage users to file change requests from within the Purify and PureCoverage graphical interface. To configure this integration between the products:</p> <ul style="list-style-type: none"> ▪ You as the installer can provide <code>rs_install</code> with site-wide default values. Note that individual users can override the default values if they have different requirements, following the instructions in the Purify and PureCoverage online help systems. (See "ClearQuest" in the online help index.) ▪ Alternatively, you can choose not to enter default values. In this case, users must enter their own values in order to use the integration. <p>The source for the following configuration values is your ClearQuest system administrator.</p>		
<p>ClearQuest client interface.</p>	<p>If you are installing on a Solaris or HP-UX platform, you can specify the default ClearQuest client for your users:</p> <ul style="list-style-type: none"> ▪ <code>x</code> for the ClearQuest native X client ▪ <code>web</code> for the ClearQuest web client <p>All other platforms support only the web client.</p> <p>Users must have ClearQuest on their path in order to use the native X client.</p>	
<p>Name of the default ClearQuest database.</p>	<p>Enter the name of a site-wide default ClearQuest database for your users who access ClearQuest through the native X client.</p> <p>Even if you have chosen <code>web</code> as your default interface, you may want to enter a value here for users who prefer to use the ClearQuest native X client.</p> <p>Source: ClearQuest system administrator.</p>	

Data for rs_install	Notes	Your Entry
URL for the ClearQuest web client.	<p>Specify the URL if you selected web as your default ClearQuest client interface.</p> <p>Even if you have chosen x as your default interface, you may want to enter a value here for users who prefer to use the web interface.</p> <p>Source: ClearQuest system administrator.</p>	

When you have gathered the information you need, go on to *Installing the products using rs_install* on page 15.

If you are installing a temporary or evaluation license

Gather the information specified in the following table:

Data for rs_install	Notes	Your Entry
The full pathname to the installation location (referred to in this chapter as Rational).	<p>This is the directory for installing all Rational Software products. You must have 20 megabytes of free disk space for each installation of Purify, Quantify, and PureCoverage. You must have 60 megabytes for a complete installation of PurifyPlus. Note that these are per-platform figures.</p> <p>The directory must be accessible from every machine on which you plan to run the Rational products—both the machines on which users <i>instrument</i> their applications, and the machines on which users <i>run</i> their applications. It must be the same for each machine, so you cannot use a local automount path like /tmp_mnt/Rational.</p> <p>If Rational does not already exist, the installation program will create it when you enter the full pathname.</p> <p>If you are installing on a read-only file system, or if you want to create this directory manually, see <i>Supplemental notes</i> on page 24. This section also shows you the structure of the directory after installation.</p>	
License key type.	Source: your Rational license key certificate or email from Rational Software.	startup or evaluation
Expiration date.	<p>Source: your Rational license key certificate or email from Rational Software.</p> <p>If you have a startup or evaluation license, enter the date in the dd-mmm-yyyy format (for example, 31-dec-2003). The field is not case sensitive.</p>	

Data for <code>rs_install</code>	Notes	Your Entry
<p>Note: If you have installed Rational ClearQuest, Rational Software's change request management system, it is possible for Purify and PureCoverage users to file change requests from within the Purify and PureCoverage graphical interface. To configure this integration between the products:</p> <ul style="list-style-type: none"> ▪ You as the installer can provide <code>rs_install</code> with site-wide default values. Note that individual users can override the default values if they have different requirements, following the instructions in the Purify and PureCoverage online help systems. (See "ClearQuest" in the online help index.) ▪ Alternatively, you can choose not to enter default values. In this case, users must enter their own values in order to use the integration. <p>The source for the following configuration values is your ClearQuest system administrator.</p>		
ClearQuest client interface.	<p>If you are installing on a Solaris or HP-UX platform, you can specify the default ClearQuest client for your users:</p> <ul style="list-style-type: none"> ▪ <code>x</code> for the ClearQuest native X client ▪ <code>web</code> for the ClearQuest web client <p>All other platforms support only the web client.</p> <p>Users must have ClearQuest on their path in order to use the native X client.</p>	
Name of the default ClearQuest database.	<p>Enter the name of a site-wide default ClearQuest database for your users who access ClearQuest through the native X client.</p> <p>Even if you have chosen <code>web</code> as your default interface, you may want to enter a value here for users who prefer to use the ClearQuest native X client.</p> <p>Source: ClearQuest system administrator.</p>	
URL for the ClearQuest web client.	<p>Specify the URL if you selected <code>web</code> as your default ClearQuest client interface.</p> <p>Even if you have chosen <code>x</code> as your default interface, you may want to enter a value here for users who prefer to use the web interface.</p> <p>Source: ClearQuest system administrator.</p>	

When you have gathered the information you need, go on to *Installing the products using `rs_install`* on page 15.

Installing the products using `rs_install`

For information about specific product and operating system versions, see the `README` file on your CD-ROM or in the directory that results when you unpack the `tar` file from Rational Software.

To install the products:

1 Make the product available for installation.

If you are installing the product from the Rational Software product CD-ROM and need instructions for mounting a CD-ROM, see page 25.

If you have obtained the products from a web or FTP download, unpack the compressed `tar` file. The directory that is created is equivalent to the top level of the CD-ROM.

2 Run the `rs_install` program. The `rs_install` program is a complete installer that guides you through whichever of the following processes are necessary:

- Setting up the license server.
- Installing product licenses.
- Installing the selected product and documentation.

Note: Your users can get online help only if you install the `html` documentation.

- Performing the post-installation tasks.

To run the `rs_install` program, go to the directory where you mounted the CD-ROM or unpacked your `tar` file. (You should not be `root` when you run `rs_install`.) For example:

```
# exit
% cd /cdrom
% ./rs_install
```

The `rs_install` program prompts you through the installation, providing detailed instructions along with default settings. The defaults appear in brackets, for example `[2]`. To accept the default, press `ENTER`.

If you select `Typical` as your installation type within `rs_install`, the entire installation, including the post-installation tasks, proceeds automatically. A typical installation installs all the products for which you have a license on the system you're logged into.

Note: After you install your license key, the `rs_install` program reminds you that you must configure your server to automatically restart the license server when it reboots. The `rs_install` program gives you instructions for doing this.

- 3 When installation is complete, go to *Step 3: Post-installation configuration tasks* on page 18 and perform any necessary post-installation procedures.

Answers to questions about `rs_install`

Below are the answers to some common questions about the `rs_install` program.

- **What's the easiest way to use `rs_install`?** Choose `Typical` when the `rs_install` program prompts you for the type of installation to perform. The program then proceeds to complete installation and post-installation tasks without further intervention from you. If you have special requirements, such as the need to install a Solaris version of the product on an HP-UX machine, or if you are setting up a new ClearQuest integration, you have to choose `Custom` instead of `Typical`; this requires you to input a small amount of additional information.
- **Can I rerun parts of the installation?** Yes. The `rs_install` program provides commands that enable you to rerun specific sections of the installation as needed. See *Using `rs_install` commands* on page 27.
- **Do I have to reenter my license server information each time I install a product?** No. You only need to enter this information once. The `rs_install` program saves the information you enter about yourself and about the machine to be used as the license server for your Rational Software product licenses in a text file called `rs_install.PurifyPlus.<version>`. The `rs_install` program reports the location of these files when you quit the program. The next time you run `rs_install`, the program uses the saved configuration information. To change your license server, use the `license_setup` command; see *Using `rs_install` commands* on page 27.

- **Do I need to install all my licenses on one server?** No. You are not required to use all of your allowed licenses for a single license server. You might want to install a product at another site and configure a license server at that site to serve the remaining licenses in your Rational Software account.
- **Which type of product license key should I install?** If you already have your permanent or TLA license key, you can install it right away. You can also request a permanent license key at www.rational.com/accountlink. Otherwise, select the startup or evaluation license to get started using the product.
Note: To ensure uninterrupted use of your Rational Software product, you should install your permanent or TLA license key as soon as possible.
- **Can I import existing user names from an earlier installation of the product installed under PureLA License Administrator with Named User licensing?** Yes. The `rs_install` script asks you if you want to import the existing `users.purela` file, and also permits you to edit the imported user names. You can also edit the user names after installation; see *Maintaining the rational.opt options file* on page 22.
- **The host I want to use as the license server for my new products is already the license server for other Rational products. How do I share the server?** You must add the new licenses to your current Rational `.dat` license file. To do this, specify the current Rational `.dat` license file as the license file name instead of using the default.
- **How do I get updates for the `rs_install` program and for the Rational products?** You can get updates from within the `rs_install` program, though you must be running the program on a machine that has internet access. The `rs_install` program's Licensing Options screen lets you select an item to download the latest version of `rs_install` (in which case `rs_install` replaces itself and restarts using the new version) or get product updates.
- **How do I report problems or make suggestions for the `rs_install` program?** You can submit comments by running the `rs_install` command with the `-report` option. This option helps you organize installation and licensing information and e-mail it to Rational technical support. You can, of course, also call or e-mail Rational technical support without this option.

Step 3: Post-installation configuration tasks

Configuration tasks include tasks that the `rs_install` program performs (or helps you perform) and tasks that you must complete manually.

The `rs_install` program performs its configuration tasks by calling the `<product>.configure` command for the product you are installing; `<product>` is `purify`, `purecov`, or `quantify`. You can also rerun the `<product>.configure` command at any time to check or to adjust your configuration; refer to *Checking and adjusting your configuration* on page 21.

The tasks that `<product>.configure` performs include:

- Configuring the cache directory.
- Setting up the online help system.
- Integrating Rational ClearQuest (for PurifyPlus, Purify, and PureCoverage).
- Running `<product>_test` to validate setup.

The tasks that you must perform manually include:

- Installing on a read-only file system; see page 18.
- Making the manual pages available; see page 19.
- Making the products available to all users; see page 19.

Note that `<producthome>` is the home directory of Purify, PureCoverage, or Quantify.

Installing on a read-only file system

Purify, PureCoverage, and Quantify work by creating and monitoring special instrumented versions of object files and libraries. They must be able to write these instrumented files to a cache directory, which by default is `Rational/releases/<producthome>/cache`.

For this reason, if you install any of the products on a file system that is mounted read-only by client machines, you must create symbolic links to a writable file system. The `rs_install` program guides you through the process of selecting a shared directory that is mounted read/write on client machines and linking the `cache` directory to this publicly writable directory.

If there is no writable shared directory mounted on client machines, have all users make a `cache` subdirectory in their home directory and set the product's `-cache-dir` option to this directory. For example:

```
% mkdir $HOME/cache
% echo $PUREOPTIONS
```

If the `PUREOPTIONS` environment variable is already set, have users specify the `-cache-dir` option:

```
csh % setenv PUREOPTIONS "-cache-dir=$HOME/cache \
    $PUREOPTIONS"
sh, ksh $ PUREOPTIONS="-cache-dir=$HOME/cache \
    $PUREOPTIONS"; export PUREOPTIONS
```

If the `PUREOPTIONS` environment variable is *not* set, have users specify:

```
csh % setenv PUREOPTIONS "-cache-dir=$HOME/cache"
sh, ksh $ PUREOPTIONS="-cache-dir=$HOME/cache"; export \
    PUREOPTIONS
```

Have all users add this same specification to their local or central `.cshrc` file, or its equivalent.

Making the manual pages available

The `rs_install` program installs the product manual pages in `Rational/releases/<producthome>/man`, where `<producthome>` is the home directory for Purify, PureCoverage, or Quantify. To make the manual pages available, do one of the following:

- Set your `MANPATH` environment variable to include `Rational/releases/<producthome>/man`.
- Copy the manual pages for the product into your `man` directory. If necessary, log in as `root` to do this.

Making the products available to all users

Note: If you are using Named User licensing, users must be listed in the `rational.opt` file in order to use Purify, PureCoverage, and Quantify; to add users to the options file, see *Maintaining the rational.opt options file* on page 22.

To make the products available to all users listed in `rational.opt`, add the full `Rational/releases/<producthome>` pathname to each user's `PATH` environment variable, or specify the full pathname in makefiles.

As an alternative to modifying your `PATH` environment variable, you can create a symbolic link to `<producthome>/<product>` from a directory such as `/usr/local/bin`. Make sure this is a symbolic link, not a copy or a hard link. Create symbolic links for each product you install, as in the following examples:

- For Purify:

```
% rm /usr/local/bin/purify
% ln -s Rational/releases/\
    <producthome>/purify /usr/local/bin
```

- For PureCoverage:

```
% rm /usr/local/bin/purecov
% ln -s Rational/releases/\
    <producthome>/purecov /usr/local/bin
```

For PureCoverage, you also need to create symbolic links to the `pc_*` script files:

```
% rm -i /usr/local/bin/pc_*
% ln -s Rational/releases/\
    <purecovhome>/scripts/pc_* /usr/local/bin
```

For more information on the `pc_*` scripts, see the PureCoverage online help system.

- For Quantify:

```
% rm /usr/local/bin/quantify
% ln -s Rational/releases/\
    <producthome>/quantify /usr/local/bin
```

For Quantify, you also need to create symbolic links to the `qv` program and to the `qx` script files:

```
% rm /usr/local/bin/qv
% rm -i /usr/local/bin/qx*
% ln -s Rational/releases/\
    <quantifyhome>/qv /usr/local/bin
% ln -s Rational/releases/\
    <quantifyhome>/qx* /usr/local/bin
```

For more information on the `qv` program and on the `qx` scripts, see the Quantify online help system.

HP-UX

- Create symbolic links for debugger scripts on HP-UX:

On HP-UX, Purify, PureCoverage, and Quantify include three scripts that enable you to start instrumented programs under a debugger. You need to create symbolic links to these scripts. For example, for Purify:

```
% rm /usr/local/bin/purify_dde
% rm /usr/local/bin/purify_xdb
% rm /usr/local/bin/purify_softdebug

% ln -s <purifyhome>/purify_dde /usr/local/bin
% ln -s <purifyhome>/purify_xdb /usr/local/bin
% ln -s <purifyhome>/purify_softdebug /usr/local/bin
```

For PureCoverage and Quantify, create the same symbolic links, substituting `purecov` or `quantify` for `purify`.

The installation is now complete. To add names to the options file, see *Maintaining the rational.opt options file* on page 22. To remove previous versions of the products, see *Removing a previous product release* on page 23.

Checking and adjusting your configuration

You can run the script `<product>.configure` (where `<product>` is `purify`, `purecov`, or `quantify`) at any time to check that your configuration is correct and to make adjustments.

To use the script, go to the product home directory. For example, for Purify, type:

```
% cd `purify -printhomedir`
```

To check your configuration and licensing, type:

```
% ./purify.configure -check
```

To run the script in interactive mode, type:

```
% ./purify.configure
```

You can also run the script in batch mode, specifying the parameters you want changed as arguments to options. For a list and description of batch mode options, type:

```
% ./purify.configure -help
```

Maintaining the rational.opt options file

Named User licensing is always used with Purify, PureCoverage, and Quantify, and is available for use with PurifyPlus. Under *Named User* licensing, the user names of all users who are authorized to run Purify, PureCoverage, and Quantify must be listed in the `rational.opt` options file. The number of user names in the file must match the number of licenses you have installed.

Users who are identified in the file can use all features of the product, including instrumenting applications, running instrumented applications, and viewing saved data files in the product's user interface. A user can run as many concurrent sessions as desired on a single host machine; this consumes a single license. The same user can run the product on additional host machines, but consumes another license for each additional machine.

The options file is created when you run the `rs_install` program. By default, this file is `Rational/config/rational.opt`. You can relocate the file yourself after installation, but you must edit the `DAEMON` line in the license file to specify the new path:

```
DAEMON rational /etc/rational /mydir/rational.opt
```

During installation, `rs_install` asks you to supply user names, one for each license you purchased. You don't have to enter all user names during installation; `rs_install` will generate dummy names to bring the total up to the number of licenses you purchased. Your entries—real names, automatically generated dummy names, or both—are recorded in the options file.

The user names are recorded in the options file in `GROUP` directives. An `INCLUDE` directive follows each `GROUP` directive, specifying one product that the users in the group are authorized to use:

```
GROUP <group name> <user1> <user2> . . . <usern>  
INCLUDE <product>:KEY=<license key> GROUP <group name>
```

For example, in the following, `alice`, `tom`, and `harry` can use Purify, but only `alice` and `harry` can use Quantify:

```
GROUP DevTools1 alice tom harry  
INCLUDE purify:KEY=456778982 GROUP DevTools1  
GROUP DevTools2 alice harry  
INCLUDE quantify:KEY=12345778654 GROUP DevTools2
```

The `KEY` is the license key from your `.dat` license file.

Modifying the list of user names

Note: If you modify the options file while the license vendor daemon is running, you must restart the license server.

You can add, change, or delete user names by running the `options_setup` script. You can also add, change, or delete user names in the options file using any text editor.

The number of users listed for each product must always match the number of licenses that you purchased. The license server must be restarted before the changes can take effect; the `options_setup` script restarts the license server for you.

For additional information about the options file, refer to your FLEXlm user's manual.

Removing a previous product release

Note: Only the installer of the product can uninstall it.

After you install the latest version of Purify, PureCoverage, or Quantify, and after all users have switched to the new version, you can remove the old release to reclaim disk space.

To remove a previous release of Purify, PureCoverage, or Quantify, go to the `Rational` directory and run the `uninstall` script:

```
% cd Rational
% config/uninstall
```

Running the `uninstall` script with no command-line arguments causes it to display the list of products in the `releases` directory. The script prompts you for the product you want to remove.

Requesting and installing the permanent or TLA license key

If you installed your product with a startup or evaluation license key, you have to enter a permanent or TLA license key before the startup or evaluation license key expires. This allows you to continue uninterrupted use of the product.

Requesting your permanent or TLA license key

To request a permanent or TLA license key, go to www.rational.com/accountlink and follow the instructions provided there.

Entering a permanent or TLA license key after initial installation

To enter your permanent or TLA license key after you have installed your Rational Software product and exited the `rs_install` program:

- 1 Go to the `Rational/releases/PurifyPlusFamily.<version>` directory and run the `license_setup` program. For instructions, see *Using rs_install commands* on page 27.
- 2 For the licensing option, select the option for setting up a permanent license.

Note: The program tells you how to update your license server machine so that it restarts the license server when it reboots. You need root permission to perform the update.

Supplemental notes

Creating an installation directory manually

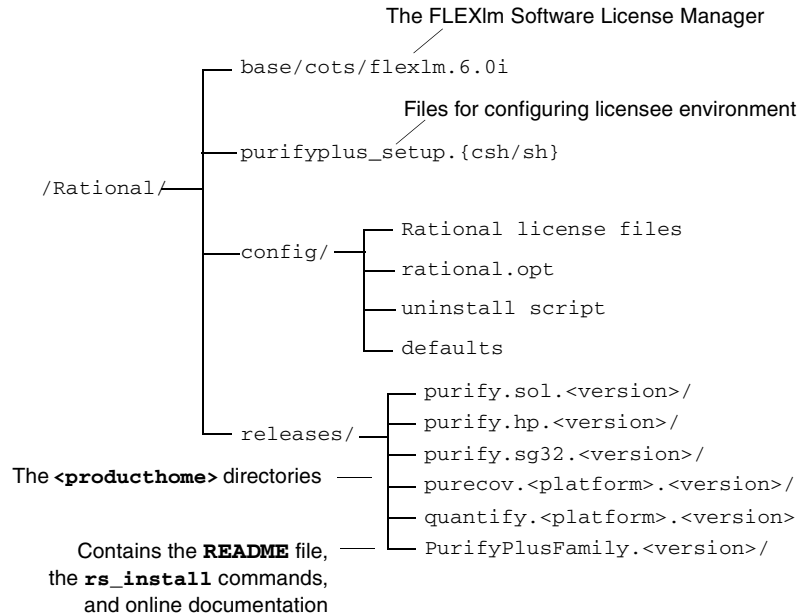
You need a publicly readable directory for the installation of your product if it is to run on multiple machines. If one does not already exist, you can create it when you run `rs_install`. You can also create it manually before you start `rs_install`.

- 1 Log into a UNIX workstation that provides access to the product files to be installed, and that mounts the file system(s) into which you want to load the product.
- 2 Create a `Rational` directory. For example:

```
% mkdir /opt/Rational
```

The `Rational` directory must be visible on all machines that are to run this product. The NFS name for `Rational` must be the same on all machines. (If you are installing the product for your use only, you can install it in your home directory.)

After the installation, the Rational directory is structured like this:




Note: Purify, PureCoverage, and Quantify must be able to write instrumented files to a cache subdirectory of the <producthome> directory. If you install on a read-only file system, you must create symbolic links to a writable file system. See *Installing on a read-only file system* on page 18.

Mounting the CD-ROM

The following instructions refer to specific operating systems. To determine your operating system, type:

```
% uname -a
```

Before you begin, make sure you know the device name of your CD-ROM drive. Ask your system administrator for this information.



On Solaris and IRIX systems with Volume Management, load the CD-ROM and then go to Step 5. (On these systems, the CD-ROM automatically mounts on the /cdrom directory. To determine whether you have Volume Management, check to see if the Solaris vold daemon or the IRIX mediad daemon is running on your system.)

To mount the CD-ROM:

1 Load the CD-ROM into the drive.

2 Log in as root:

```
% su root
```

3 If you do not already have one, create a `cdrom` directory to be the mount point for the CD-ROM drive:

```
# mkdir /cdrom
```

4 Mount the CD-ROM:



▫ On Solaris systems without Volume Management:

```
# /etc/mount -r -F hsfs <cdrom-device-name> /cdrom
```



▫ If your HP-UX system is configured to mount the CD-ROM at `/cdrom`:

```
# /etc/mount /cdrom
```

▫ If your HP-UX system is not configured to mount the CD-ROM at `/cdrom`, use the following command:

```
# /etc/mount -r -F cdfs <cdrom-device-name> /cdrom
```



▫ On IRIX 6.x:

```
# /etc/mount -r -t iso9660 <cdrom-device-name> /CDROM
```

5 To verify that the CD-ROM is mounted, use the `ls` command to list the files:

```
# ls -R /cdrom
```

Ejecting the CD-ROM

After you complete the installation, eject the CD-ROM.



On Solaris with Volume Management, type:

```
% eject cdrom
```

On Solaris without Volume Management, type:

```
% su root  
# umount /cdrom  
# eject cdrom  
# exit
```



On HP-UX, type:

```
% su root
# umount /cdrom
# exit
```

Press the eject button on the CD-ROM drive.



On IRIX, type:

```
% eject /CDROM
```

Using rs_install commands

The `rs_install` program includes four commands that you can use to rerun specific sections of the `rs_install` program without actually reinstalling any products: `license_setup`, `license_check`, `post_install`, and `options_setup`.

To use these commands, go to the `PurifyPlusFamily.<version>` directory. For example:

```
% cd Rational/releases/PurifyPlusFamily.<version>
% ./license_setup
```

Use these commands as follows:

- Use the `license_setup` command to rerun the license setup phase of the installation. Use `license_setup` to import your permanent or TLA license keys and whenever you want to change your licensing information.
- Use the `license_check` command to check your license server and the `.dat` license file to make sure your license information is correct.
- Use the `post_install` command to rerun the post-installation phase of the installation. One of the actions this command performs is to call the `<product>.configure` command; see *Step 3: Post-installation configuration tasks* on page 18.
- Use `options_setup` to modify the list of users allowed to use the Rational Software product. For more information, see *Modifying the list of user names* on page 23.

Using the FLEXlm Software License Manager

The FLEXlm Software License Manager monitors license access, simultaneous usage, idle time, and so on. It includes the following components:

- A vendor daemon named `rational` that dispenses Purify, PureCoverage, and Quantify licenses. The `rational` daemon is used for all licensed Rational Software products. If you have products from other vendors that also use FLEXlm, they will include their own vendor daemons.
- A license manager daemon named `lmgrd` that is used by all licensed products from all vendors that use FLEXlm. The `lmgrd` daemon does not process requests on its own, but forwards requests to the appropriate vendor daemon.
- A Rational license file that specifies your license servers, vendor daemons, and product licenses.

The Rational license file

The Rational `.dat` license file is a text file that in most cases is created when you run the `rs_install` or `license_setup` program. The `.dat` license file is based on data from the `.upd` file that you receive from AccountLink.

The file for startup and evaluation licenses is:

```
Rational/config/Temporary.dat
```

The default file for permanent or TLA licenses is:

```
Rational/config/<server name>.dat
```

Note: For best results, use the Rational license file only for Rational Software product licenses.

The `rs_install` program saves the license path to `<producthome>/.lm_license_file`. This is the path that Purify, PureCoverage, and Quantify use to locate the license file. You can override the location in `.lm_license_file` by setting the `LM_LICENSE_FILE` environment variable. The full path searched is equivalent to `$LM_LICENSE_FILE: 'cat.lm_license_file'`.

Verifying that FLEXlm is working

To verify that your FLEXlm License Manager is operational and that the daemons are running, type the following on your license server:

```
% /bin/ps -e | egrep "lmgrd|rational"
```

The output should include lines similar to the following:

```
/bin/ps -e | egrep "lmgrd|rational"  
 351 ? 0:00 rational  
 345 ? 0:01 lmgrd
```

Using FLEXlm commands

The FLEXlm License Manager supports the following commands for system administration:

Use this command:	To:
lmdiag	Diagnose problems when you cannot check out a license
lmdown	Shut down the license and vendor daemons
lmhostid	Report the license manager host ID of a workstation
lmreread	Reread the license file and start new vendor daemons
lmstat	Report status on daemons and feature usage
exinstal	Report on licenses in the license file you specify on the command line

Learning more about FLEXlm

For more information about the FLEXlm Software License Manager, see the *FLEXlm End User Manual* that is included on your Rational Software CD-ROM.

The *FLEXlm End User Manual*, along with answers to frequently asked questions about FLEXlm, is also available at <http://www.globetrotter.com/manual.htm>.

Rational Purify: What it does

Rational® Purify® is the most comprehensive runtime error detection tool available. It checks all the code in your program, including any application, system, and third-party libraries. Purify works with complex software applications, including multi-threaded and multi-process applications.

Purify checks every memory access operation, pinpointing *where* errors occur and providing detailed diagnostic information to help you analyze *why* the errors occur. Among the many errors that Purify helps you locate and understand are:

- Reading or writing beyond the bounds of an array
- Using uninitialized memory
- Reading or writing freed memory
- Reading or writing beyond the stack pointer
- Reading or writing through null pointers
- Leaking memory and file descriptors

With Purify, you can develop clean code from the start, rather than spending valuable time debugging problem code later.

This chapter introduces the basic concepts involved in using Purify. For complete information, see the Purify online help system.

Finding errors in Hello World

This chapter shows you how to use Purify to find memory errors in an example Hello World program. If you run the example yourself, you should expect minor platform-related differences in program output from what is shown here.

Before you begin:

- 1 Create a new working directory. Go to the new directory and copy the `hello_world.c` program and related files from the `<purifyhome>/example` directory. For example:

```
% mkdir /usr/home/chris/pwork
% cd /usr/home/chris/pwork
% cp <purifyhome>/example/hello* .
```

- 2 Examine the code in `hello_world.c`. The version of `hello_world.c` provided with Purify is slightly different from the traditional version.

```
1 /*
2  * Copyright (c) 1992-1997 Rational Software Corp.
3  * ...
9  * This is a test program used in Purifying Hello World
10 */
11
12 #include <stdio.h>
13 #include <malloc.h>
14
15 static char *helloWorld = "Hello, World";
16
17 main()
18 {
19     char *mystr = malloc(strlen(helloWorld));
20
21     strncpy(mystr, helloWorld, 12);
22     printf("%s\n", mystr);
23 }
```

At first glance there are no obvious errors, yet the program actually contains a memory access error and leaked memory that Purify will help you to identify.

Instrumenting a program

- 1 Compile and link the Hello World program, then run the program to verify that it produces the expected output:

```
% cc -g hello_world.c
% a.out
```

output ——— Hello, World

- 2 Instrument the program by adding `purify` to the front of the compile/link command line. To get the maximum amount of detail in Purify messages, use the `-g` option:

```
% purify cc -g hello_world.c
```



Note: On IRIX, you can add `purify` in front of the compile/link command line, or you can Purify the executable:

```
% purify a.out
```

You then run the instrumented program by typing:

```
% a.out.pure
```



On IRIX, Purify caches Dynamic Shared Objects (DSOs), not object files. References to linkers and link-line options in this book do not apply to Purify on IRIX.

Compiling and linking in separate stages

If you compile and link your program in separate stages, specify `purify` only on the link line. For example:

On the compile line, use:

```
% cc -c -g hello_world.c
```

On the link line, use:

```
% purify cc -g hello_world.o
```

Running the instrumented program

Run the instrumented Hello World program:



```
% a.out
```



On IRIX, if you use `purify` on the executable instead of on the compile/link line, type:

```
% a.out.pure
```

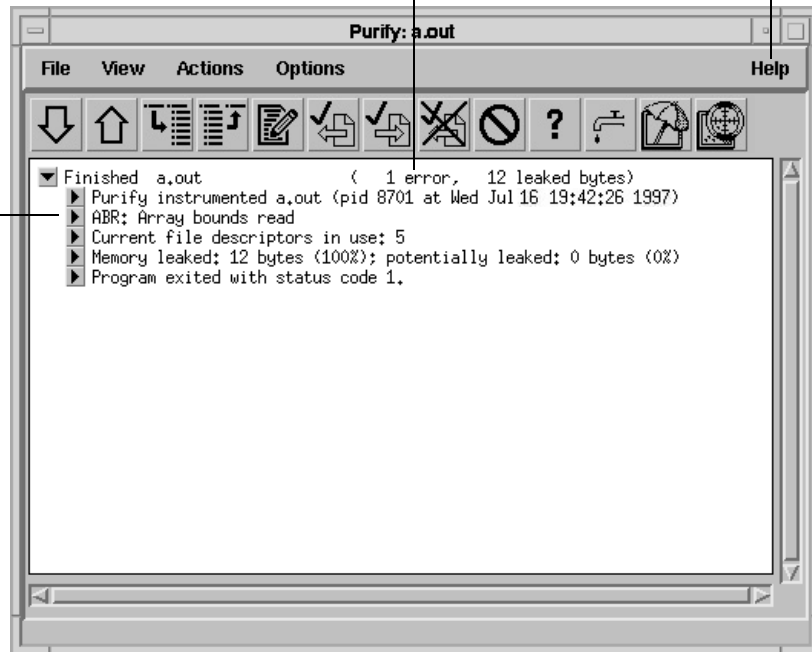
This prints “Hello, World” in the current window and displays the Purify Viewer.

Purify displays the number of access errors and leaked bytes detected

Click for a list of Purify error messages

The Purify Viewer displays messages about the program, including errors such as this ABR error

For a description of a message, right click the message, then select **Explain message** from the pop-up menu

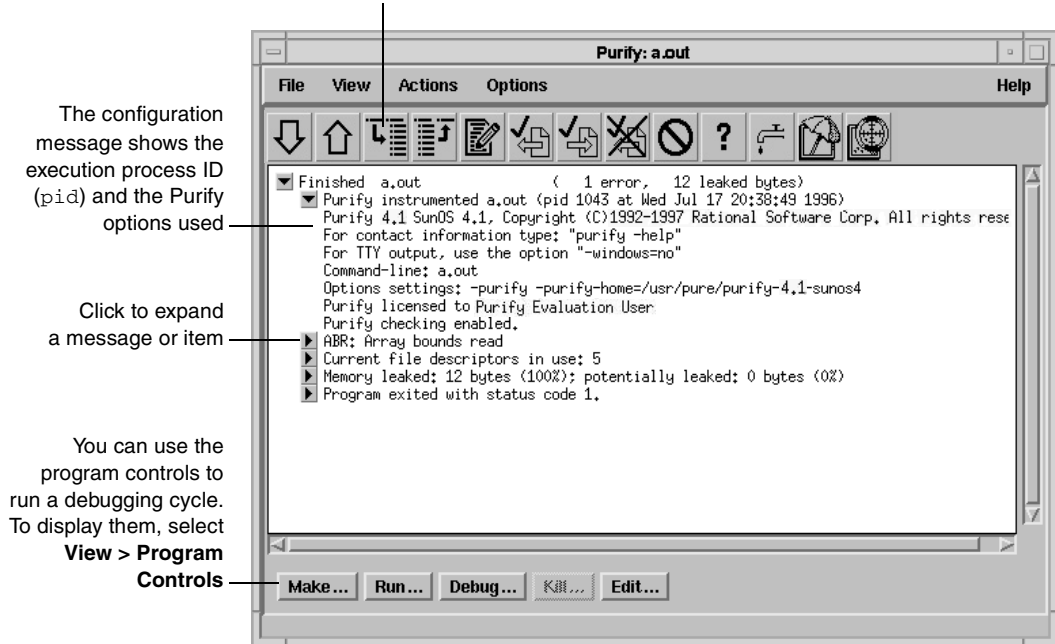


Notice that the instrumented Hello World program starts, runs, and exits normally. Purify does not stop the program when it finds an error.

Seeing all your errors at a glance

The Purify Viewer displays the results of the run of the instrumented Hello World program. You can expand each message to see additional details.

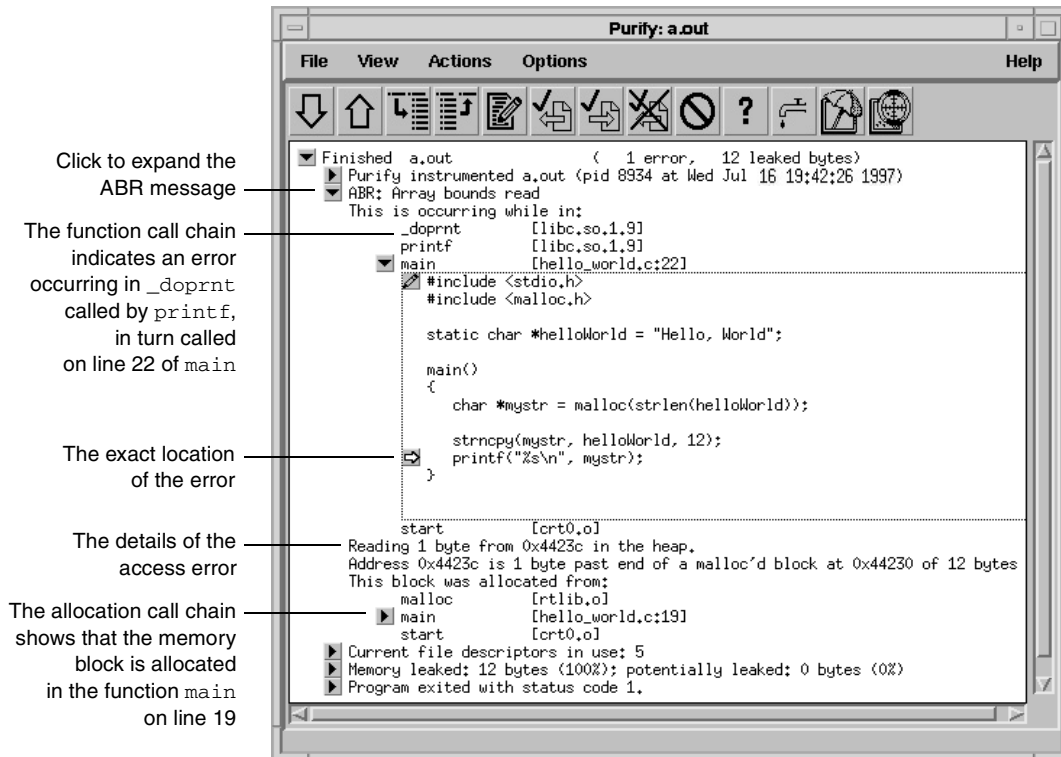
Select one or more messages in the Viewer, then click to expand the messages



Note: The Viewer displays messages for a single executable only. It is specific to the name of the executable, the directory containing the executable, and the user ID.

Finding and correcting errors

Purify reports an array bounds read (ABR) memory access error in the Hello World program. You can expand the ABR message to see the exact location of the error.



Note: To make debugging easier, Purify reports line numbers, source filenames, and local variable names whenever possible if you use the `-g` compiler option when you instrument the program. If you do not use the `-g` option, Purify reports only function names and object filenames.



On IRIX, system libraries retain their source file and line number information; therefore, the ► can appear next to a system library function whose source file is not available. When you click the ► for such a line, Purify prompts you for the location of the source file. Enter the location of the file if you know it, and then click **OK** to expand the line.


Understanding the cause of the error

To understand the cause of the ABR error, look at the code in `hello_world.c` again.

```
.
.
.
15 static char *helloWorld = "Hello, World";
16
17 main()
18 {
19     char *mystr = malloc(strlen(helloWorld));
20
21     strncpy(mystr, helloWorld, 12);
22     printf("%s\n", mystr);
23 }
```

Purify reports that the ABR error occurs here

On line 22, the program requests `printf` to display `mystr`, which is initialized by `strncpy` on line 21 for the 12 characters in “Hello, World.” However, `_doprnt` is accessing one byte more than it should. It is looking for a `NULL` byte to terminate the string. The extra byte for the string’s `NULL` terminating character has *not* been allocated and initialized.

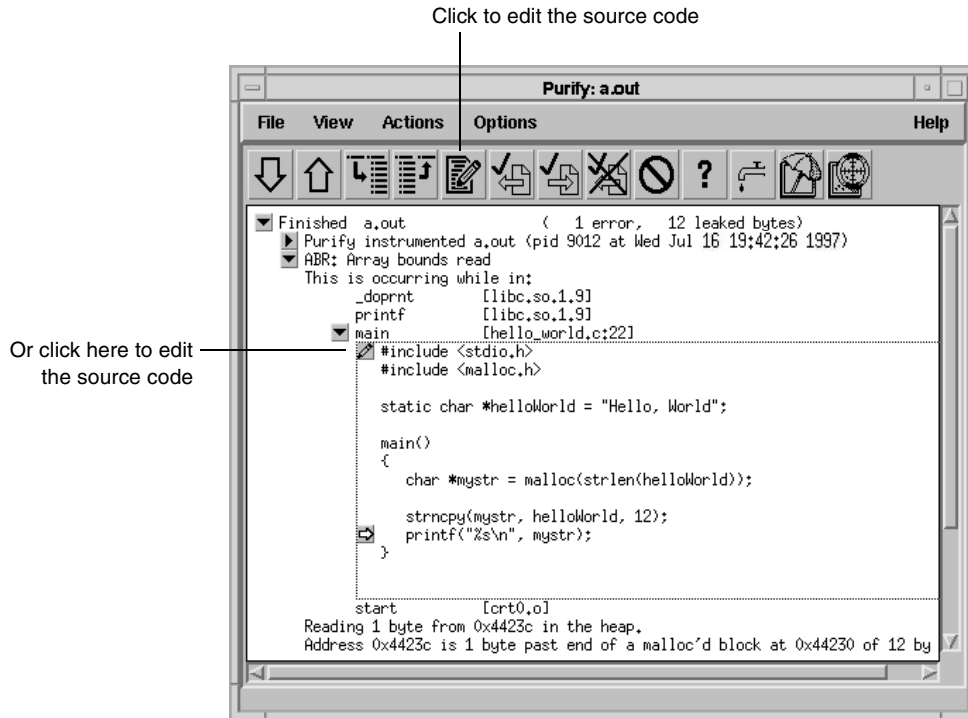


For more information, see *How Purify finds memory-access errors* on page 50.

Correcting the ABR error

To correct this ABR error:

- 1 Click the Edit tool  to open an editor.



Note: By default, Purify displays seven lines of the source code file in the Viewer. You can change the number of lines of source code displayed by setting an X resource.

- 2 Change lines 19 and 21 as follows:

```
19 char *mystr = malloc(strlen(helloWorld)+1);
20
21 strncpy(mystr, helloWorld, 13);
```


Finding leaked memory

When a program exits, Purify searches for memory leaks and reports all memory blocks that were allocated but for which no pointers exist.

Note: When you run longer-running instrumented programs, you can click the New Leaks tool to generate a new leaks summary while the program is running.

- 1 Expand the memory-leaked summary for Hello World.

The memory-leaked summary shows the number of leaked bytes as a percentage of the total heap size. If there is more than one memory leak, Purify sorts them by the number of leaked bytes, displaying the largest leaks first.

- 2 Expand the MLK message.

When you run your programs, click the New Leaks tool to generate a new leaks summary while the program is running

The memory-leaked summary reports 12 bytes of leaked memory

The call chain shows how the leaked memory was allocated

Memory analysis by category

	Blocks	Bytes
Leaked	1	12
Potentially Leaked	0	0
In-Use	0	0

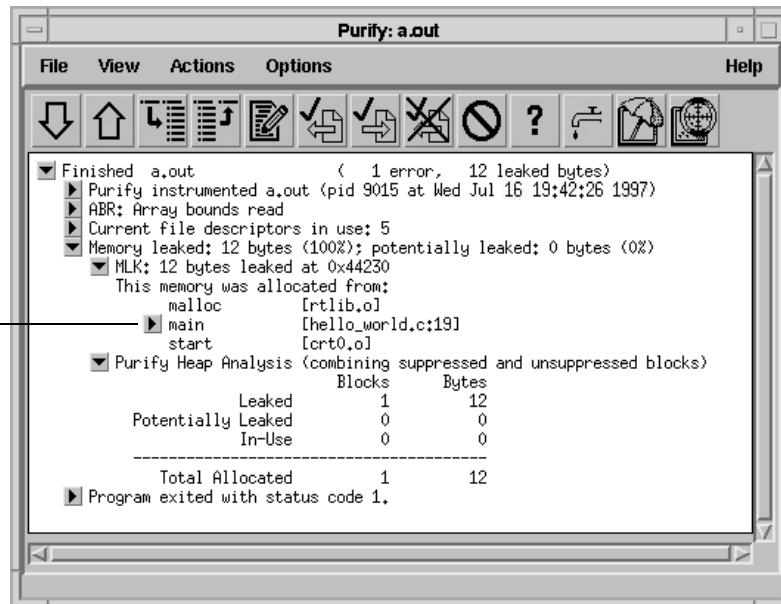
Total Allocated	1	12

```
Finished a.out ( 1 error, 12 leaked bytes)
└─ Purify instrumented a.out (pid 9015 at Wed Jul 16 19:42:26 1997)
└─ ABR: Array bounds read
└─ Current file descriptors in use: 5
└─ Memory leaked: 12 bytes (100%); potentially leaked: 0 bytes (0%)
  └─ MLK: 12 bytes leaked at 0x44230
    This memory was allocated from:
      malloc [rtlib.o]
      └─ main [hello_world.c:191]
        start [crt0.o]
  └─ Purify Heap Analysis (combining suppressed and unsuppressed blocks)
    ┌─ Leaked 1 12
    ┌─ Potentially Leaked 0 0
    ┌─ In-Use 0 0
    └─ -----
    Total Allocated 1 12
  └─ Program exited with status code 1.
```

Correcting the MLK error


It is not immediately obvious why this memory was leaked. If you look closer, however, you can see that this program does not have an `exit` statement at the end. Because of this omission, the `main` function returns rather than calls `exit`, thereby making `mystr`—the only reference to the allocated memory—go out of scope.

Line 19 of `hello_world.c` in `main` allocates 12 bytes of leaked memory. The start of this memory block is `0x44230`, the same block with the array bounds read error in `_doprnt`



If `main` called `exit` at the end, `mystr` would remain in scope at program termination, retaining a valid pointer to the start of the allocated memory block. Purify would then have reported it as memory in use rather than memory leaked. Alternatively, `main` could `free mystr` before returning, deallocating the memory so it is no longer in use or leaked.

To correct this MLK error:

- 1 Click the Edit tool  to open an editor.
- 2 Add a call to `exit(0)` at the end of the Hello World program.

Looking at the heap analysis

Purify distinguishes between three memory states, reporting both the number of blocks in each state and the sum of their sizes:

- Leaked memory
- Potentially leaked memory
- Memory in use

A true memory leak (MLK) is memory to which your program has no pointer

A potential memory leak (PLK) is memory that does not have a pointer to its beginning, but does have one to its interior

Memory in use (MIU) is memory to which your program has pointers (these are not leaks)

```
Finished a.out ( 1 error, 12 leaked bytes)
└─ Purify instrumented a.out (pid 9015 at Wed Jul 16 19:42:26 1997)
└─ ABR: Array bounds read
└─ Current file descriptors in use: 5
└─ Memory leaked: 12 bytes (100%); potentially leaked: 0 bytes (0%)
  └─ MLK: 12 bytes leaked at 0x44230
    This memory was allocated from:
      malloc [rtlib.o]
      main [hello_world.c:19]
      start [crt0.o]
  └─ Purify Heap Analysis (combining suppressed and unsuppressed blocks)
    Blocks Bytes
    Leaked 1 12
    Potentially Leaked 0 0
    In-Use 0 0
    -----
    Total Allocated 1 12
  └─ Program exited with status code 1.
```

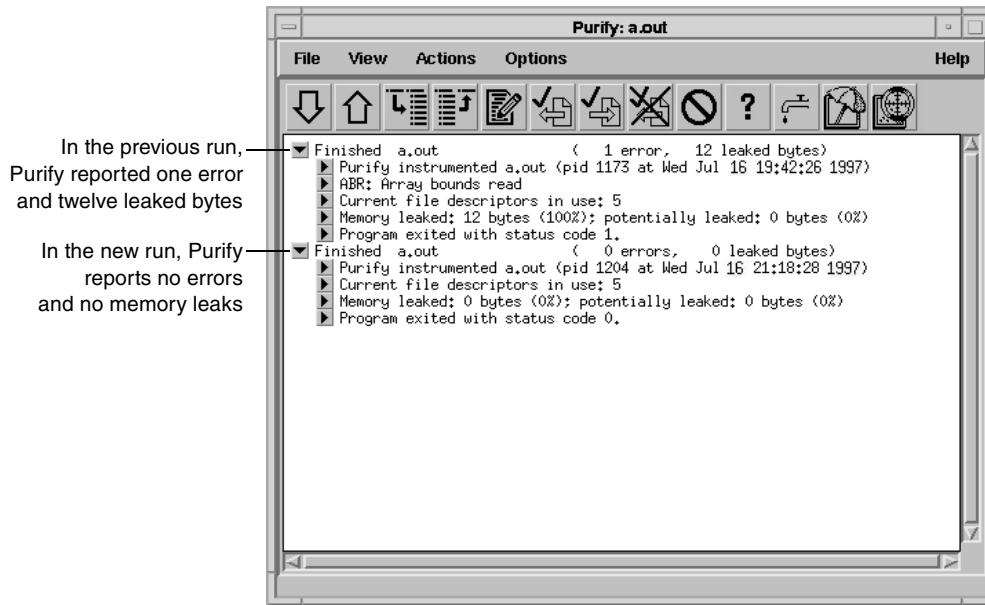
The exit status message provides information about:

- *Basic memory usage containing* statistics not easily available from a single shell command. It includes program code and data size, as well as maximum heap and stack memory usage in bytes.
- *Shared-library memory usage* indicating which libraries were dynamically linked and their sizes.

Comparing program runs

To verify that you have corrected the ABR and MLK errors, recompile the program with `purify`, and run it again.

Purify displays the results of the new run in the same Viewer as the previous run so it's easy to compare them. In this simple Hello World program, you can quickly see that the new run no longer contains the ABR and MLK errors.



Congratulations! You have successfully Purify'd the Hello World program.

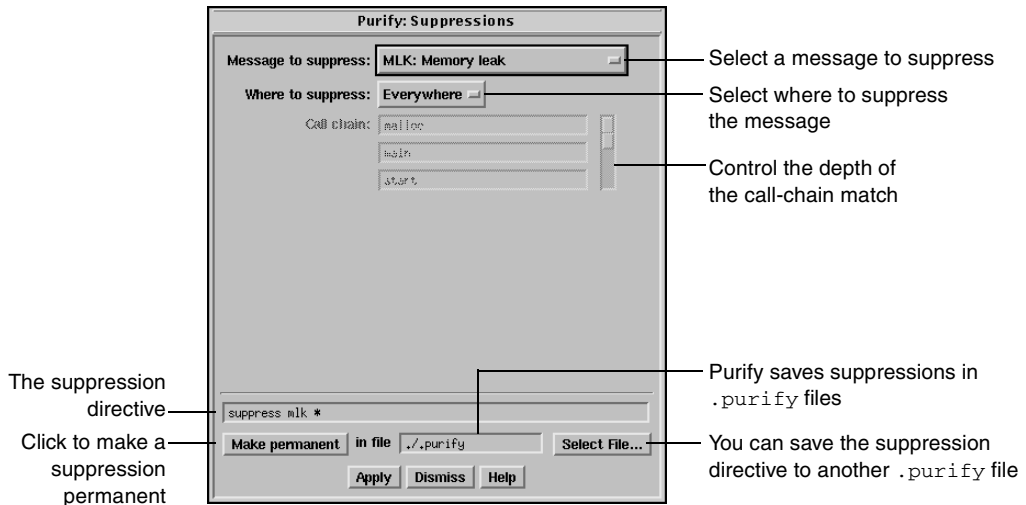
Suppressing Purify messages

A large program can generate hundreds of error messages. To quickly focus on the most critical ones, you can suppress the less critical messages based on their type and source. For example, you might want to hide all informational messages, or hide all messages that originate in a specific file.

You can suppress messages in the Viewer either during or after a run of your program. To suppress a message in the Viewer:

- 1 Select the message you want to suppress.
- 2 Select **Options > Suppressions**.

Purify displays the Suppressions dialog, containing information about the selected message.



You can also specify suppressions directly in a `.purify` file. Suppressions created in the Viewer take precedence over suppressions in `.purify` files; however, they apply only to the current Purify session. Unless you click **Make permanent**, they do not remain when you restart the Viewer.

Saving Purify output to a view file

A view file is a binary representation of all messages generated in a Purify run that you can browse with the Viewer or use to generate reports independent of a Purify run. You can save a run to a view file to compare the results of one run with the results of subsequent runs, or to share the file with other developers.

Saving a run to a view file from the Viewer

To save a program run to a view file from the Viewer:

- 1 Wait until the program finishes running, then click the run to select it.
- 2 Select **File > Save As**.
- 3 Type a filename, using the `.pv` extension to identify the run as a Purify view file.

Opening a view file

To open a view file from the Viewer:

- 1 Select **File > Open**.
- 2 Select the view file you want to open.

Purify displays the run from the view file in the Viewer. You can work with the run just as you would if you had run the program from the Viewer.

You can also use the `-view` option to open a view file. For example:

```
% purify -view <filename>.pv
```

This opens the `<filename>.pv` view file in a new Viewer.

Using your debugger with Purify

You can run an instrumented program directly under your debugger so that when Purify finds an error, you can investigate it immediately.

Alternatively, you can enable Purify's just-in-time (JIT) debugging feature to have Purify start your debugger *only* when it encounters an error—and you can specify which types of errors trigger the debugger. JIT debugging is useful for errors that appear only once in a while. When you enable JIT debugging, Purify suspends execution of your program just before the error occurs, making it easier to analyze the error.

Using Purify with PureCoverage



Purify is designed to work closely with PureCoverage, Rational Software's runtime test coverage tool. PureCoverage identifies the parts of your program that have not yet been tested so you can tell whether you're exercising your program sufficiently for Purify to find all the memory errors in your code.

To use Purify with PureCoverage, add both product names to the front of your link line. Include all ' with the program to which they refer. For example:

```
% purify <purifyoptions> purecov <purecovoptions> \  
cc -g hello_world.c -o hello_world
```

To start PureCoverage from the Purify Viewer, click the PureCoverage icon  in the toolbar.

For more information, see *Rational Purify: What it does* on page 31.

Purify API functions

You can call Purify's API functions from your source code or from your debugger to gain more control over Purify's error checking. By calling these functions from your debugger, you get additional control without modifying your source code. You can use Purify's API functions to check memory state and to search for memory and file-descriptor leaks.

For example, by default Purify reports memory leaks only when you exit your program. However, if you call the API function `purify_new_leaks` at key points throughout your program, Purify reports the memory leaks that have occurred since the last time the function was called. This periodic checking enables you to locate and track memory leaks more effectively.

To use Purify API functions, include `<purifyhome>/purify.h` in your code and link with `<purifyhome>/purify_stubs.a`.

Commonly used functions	Description
<code>int purify_describe (char *addr)</code>	Prints specific details about memory
<code>int purify_is_running (void)</code>	Returns "TRUE" if the program is instrumented

Commonly used functions	Description
<code>int purify_new_inuse (void)</code>	Prints a message on all memory newly in use
<code>int purify_new_leaks (void)</code>	Prints a message on all new leaks
<code>int purify_new_fds_inuse (void)</code>	Lists the new open file descriptors
<code>int purify_printf (char *format, ...)</code>	Prints formatted text to the Viewer or log-file
<code>int purify_watch (char *addr)</code>	Watches for memory write, malloc, free
<code>int purify_watch_n (char *addr, int size, char *type)</code>	Watches memory: type = "r", "w", "rw"
<code>int purify_watch_info (void)</code>	Lists active watchpoints
<code>int purify_watch_remove (int watchno)</code>	Removes a specified watchpoint
<code>int purify_what_colors (char *addr, int size)</code>	Prints the color coding of memory

Build-time options

Specify build-time options on the link line when you instrument a program with Purify. For example:

```
% purify -cache-dir=$HOME/cache -always-use-cache-dir cc ...
```

Commonly used build-time options	Default
<code>-always-use-cache-dir</code> Forces all instrumented object files to be written to the global cache directory	no
<code>-cache-dir</code> Specifies the global directory where Purify caches instrumented object files	<purifyhome>/cache
<code>-collector</code> Specifies the collect program to handle static constructors (for use with gcc, g++)	none
<code>-ignore-runtime-environment</code> Prevents the runtime Purify environment from overriding the option values used in building the program	no

Commonly used build-time options	Default
-linker Sets the alternative linker to build the executables instead of the system default	system-dependent
-print-home-dir Prints the name of the directory where Purify is installed, then exits	

Conversion characters for filenames

Use these conversion characters when specifying filenames for options such as `-log-file` and `-view-file`.

Character	Converts to
%V	Full pathname of program with "/" replaced by "_"
%v	Program name
%p	Process id (pid)
qualified filenames (./%v.pv)	Absolute or relative to current working directory
unqualified filenames (no '/')	Directory containing the program

Runtime options

Specify runtime options on the link line or by using the `PURIFYOPTIONS` environment variable. For example:

```
% setenv PURIFYOPTIONS "-log-file=mylog.%v.%p 'printenv PURIFYOPTIONS'"
```

Commonly used runtime options	Default
-auto-mount-prefix Removes the prefix used by file system auto-mounters	/tmp_mnt
-chain-length Sets the maximum number of stack frames to print in a report	6
-fds-in-use-at-exit Specifies that the file descriptor in use message be displayed at program exit	yes

Commonly used runtime options	Default
-follow-child-processes Controls whether Purify monitors child processes in an instrumented program	no
-jit-debug Enables just-in-time debugging	none
-leaks-at-exit Reports all leaked memory at program exit	yes
† -log-file Writes Purify output to a log file instead of the <i>Viewer</i> window	stderr
-messages Controls display of repeated messages: "first", "all", or in a "batch" at program exit	first
-program-name Specifies the full pathname of the instrumented program if <code>argv[0]</code> contains an undesirable or incorrect value	<code>argv[0]</code>
-show-directory Shows the directory path for each file in the call chain, if the information is available	no
-show-pc Shows the full pc value in each frame of the call chain	no
-show-pc-offset Appends a pc-offset to each function name in the call chain	no
† -view-file Saves Purify output to a view file (<code>.pv</code>) instead of the <i>Viewer</i> .	none
-user-path Specifies a list of directories in which to search for programs and source code	none
-windows Redirects Purify output to <code>stderr</code> instead of the <i>Viewer</i> if <code>-windows=no</code>	none

† Can use the conversion characters listed on page 47.

Purify messages

Purify reports the following messages. For detailed, platform-specific information, see the Purify online help system.

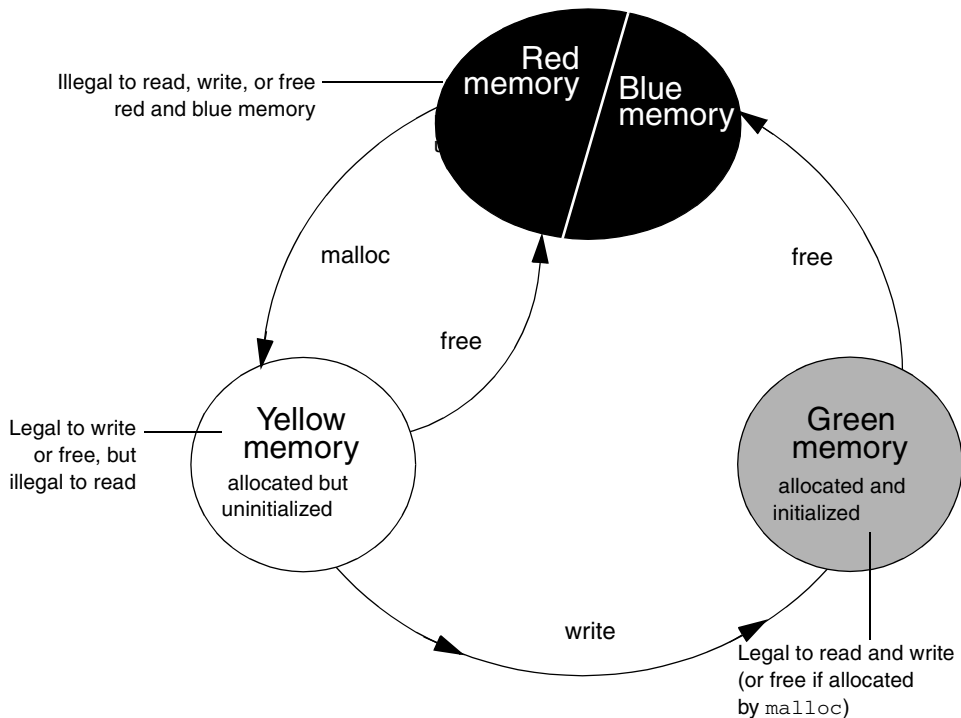
Message	Description	Severity*	Message	Description	Severity*
ABR	Array Bounds Read	W	NPR	Null Pointer Read	F
ABW	Array Bounds Write	C	NPW	Null Pointer Write	F
BRK	Misuse of Brk or Sbrk	C	PAR	Bad Parameter	W
BSR	Beyond Stack Read	W	PLK	Potential Leak	W
BSW	Beyond Stack Write	W	PMR	Partial UMR	W
COR	Core Dump Imminent	F	SBR	Stack Array Bounds Read	W
FIU	File Descriptors In Use	I	SBW	Stack Array Bounds Write	C
FMM	Freeing Mismatched Memory	C	SIG	Signal	I
FMR	Free Memory Read	W	SOF	Stack Overflow	W
FMW	Free Memory Write	C	UMC	Uninitialized Memory Copy	W
FNH	Freeing Non Heap Memory	C	UMR	Uninitialized Memory Read	W
FUM	Freeing Unallocated Memory	C	WPF	Watchpoint Free	I
IPR	Invalid Pointer Read	F	WPM	Watchpoint Malloc	I
IPW	Invalid Pointer Write	F	WPN	Watchpoint Entry	I
MAF	Malloc Failure	I	WPR	Watchpoint Read	I
MIU	Memory In-Use	I	WPW	Watchpoint Write	I
MLK	Memory Leak	W	WPX	Watchpoint Exit	I
MRE	Malloc Reentrancy Error	C	ZPR	Zero Page Read	F
MSE	Memory Segment Error	W	ZPW	Zero Page Write	F

* Message severity: F=Fatal, C=Corrupting, W=Warning, I=Informational

How Purify finds memory-access errors

Purify monitors every memory operation in your program, determining whether it is legal. It keeps track of memory that is not allocated to your program, memory that is allocated but uninitialized, memory that is both allocated and initialized, and memory that has been freed after use but is still initialized.

Purify maintains a table to track the status of each byte of memory used by your program. The table contains two bits that represent each byte of memory. The first bit records whether the corresponding byte has been allocated. The second bit records whether the memory has been initialized. Purify uses these two bits to describe four states of memory: red, yellow, green, and blue.



Purify checks each memory operation against the color state of the memory block to determine whether the operation is valid. If the program accesses memory illegally, Purify reports an error.

- *Red*: Purify labels heap memory and stack memory red initially. This memory is unallocated and uninitialized. Either it has never been allocated, or it has been allocated and subsequently freed.

In addition, Purify inserts guard zones around each allocated block and each statically allocated data item, in order to detect array bounds errors. Purify colors these guard zones red and refers to them as *red zones*. It is illegal to read, write, or free red memory because it is not owned by the program.

- *Yellow*: Memory returned by `malloc` or `new` is yellow. This memory has been allocated, so the program owns it, but it is uninitialized. You can write yellow memory, or free it if it is allocated by `malloc`, but it is illegal to read it because it is uninitialized. Purify sets stack frames to yellow on function entry.
- *Green*: When you write to yellow memory, Purify labels it green. This means that the memory is allocated and initialized. It is legal to read or write green memory, or free it if it was allocated by `malloc` or `new`. Purify initializes the *data* and *bss* sections of memory to green.
- *Blue*: When you free memory after it is initialized and used, Purify labels it blue. This means that the memory is initialized, but is no longer valid for access. It is illegal to read, write, or free blue memory.

Since Purify keeps track of memory at the byte level, it catches all memory-access errors. For example, it reports an uninitialized memory read (UMR) if an `int` or `long` (4 bytes) is read from a location previously initialized by storing a `short` (2 bytes).

How Purify checks statically allocated memory

In addition to detecting access errors in dynamic memory, Purify detects references beyond the boundaries of data in global variables and static variables; that is, data allocated statically at link time as opposed to dynamically at run time.

Here is an example of data that is handled by the static checking feature:

```
int array[10];
main() {
    array[11] = 1;
}
```

In this example, Purify reports an array bounds write (ABW) error at the assignment to `array[11]` because it is 4 bytes beyond the end of the array.

Purify inserts red zones around each variable in your program's static-data area. If the program attempts to read from or write to one of these red zones, Purify reports an array bounds error (ABR or ABW).

Purify inserts red zones into the data section *only* if all data references are to known data variables. If Purify finds a data reference that is relative to the start of the data section as opposed to a known data variable, Purify is unable to determine which variable the reference involves. In this case, Purify inserts red zones at the beginning and end of the data section only, not between data variables.

Purify provides several command-line options and directives to aid in maximizing the benefits of static checking.

Rational PureCoverage: What it does

During the development process, software changes daily, sometimes hourly. Unfortunately, test suites do not always keep pace. Rational® PureCoverage® is a simple, easily deployed tool that identifies the lines and functions in your code that have not been exercised by testing. It supports C and C++ applications, as well as Java applications running on a Solaris SPARC 32-bit Java virtual machine (JVM).

Using PureCoverage, you can:

- Pinpoint untested areas of your code
- Accumulate coverage data over multiple runs and multiple builds
- Merge data from different programs sharing common source code
- Work closely with Purify to make sure that Purify finds errors throughout your *entire* application
- Automatically generate a wide variety of useful reports
- Access the coverage data so you can write your own reports
- Collect coverage data on UNIX for viewing on a Windows system

PureCoverage provides the information you need to identify gaps in testing quickly, saving time and effort.

This chapter introduces the basic concepts involved in using PureCoverage. For complete information, see the PureCoverage online help system, including the *Java Supplement* for PureCoverage.

Finding untested Java code

PureCoverage provides accurate coverage information that identifies all the gaps in your testing of Java code. To collect Java code coverage data, run PureCoverage with the `-java` option, as follows:

- For an applet:

```
% purecov [<PureCoverage options>] -java \  
<applet viewer> [<applet viewer options>] <html file>
```

- For a class file:

```
% purecov [<PureCoverage options>] -java \  
<Java executable> [<Java options>] <class>
```

- For a JAR file:

```
% purecov [<PureCoverage options>] -java \  
<Java executable> [<Java options>] <JAR switch> \  
<JAR file>.jar
```

- For a container program:

```
% purecov [<PureCoverage options>] -java <exename> \  
[<arguments to exename>]
```

PureCoverage also monitors native methods written in languages such as C/C++ and assembly using the Java Native Interface (JNI).

To display the coverage data for the program, use the command:

```
% purecov -view java.pcv
```

For an example showing how to use PureCoverage to monitor Java code, and for information about ways to control code monitoring, see the *Java Supplement* for PureCoverage, which is included with the PureCoverage online help system.

Finding untested C/C++ code

This chapter shows you how to use PureCoverage to find the untested parts of the `hello_world.c` program.

Before you begin:

- 1 Create a new working directory. Go to the new directory, and copy the `hello_world.c` program and related files from the `<purecovhome>/example` directory:

```
% mkdir /usr/home/pat/example
% cd /usr/home/pat/example
% cp <purecovhome>/example/hello* .
```

- 2 Examine the code in `hello_world.c`.

The version of `hello_world.c` provided with PureCoverage is slightly more complicated than the usual textbook version.

```
#include <stdio.h>
void display_hello_world();
void display_message();

main(argc, argv)
int argc;
char** argv;
{
    if (argc == 1)
        display_hello_world();
    else
        display_message(argv[1]);
    exit(0);
}

void
display_hello_world()
{
    printf("Hello, World\n");
}

void
display_message(s)
    char *s;
{
    printf("%s, World\n", s);
}
```

Instrumenting a C/C++ program

- 1 Compile and link the Hello World program, then run the program to verify that it produces the expected output:

```
% cc -g hello_world.c
% a.out
```

output ————— Hello, World

- 2 Instrument the program by adding `purecov` to the front of the compile/link command line. To have PureCoverage report the maximum amount of detail, use the `-g` option:

```
% purecov cc -g hello_world.c
```

Note: If you compile your code *without* the `-g` option, PureCoverage provides only function-level data. It does not show line-level data.

A message appears, indicating the version of PureCoverage that is instrumenting the program:

```
PureCoverage 2003.06.00 Solaris 2 (32-bit), Copyright
(C)1994-2003 Rational Software Corp.
All rights reserved.
Instrumenting: hello_world.o Linking
```

Note: When you compile and link in separate stages, add `purecov` only to the link line.

Running the instrumented C/C++ program

Run the instrumented Hello World program:

```
% a.out
```

PureCoverage displays the following:

	Name of the instrumented executable	You can use this command to display technical support contact information
Start-up banner	**** PureCoverage instrumented a.out (pid 3466 at Wed Feb 6 10:32:40 2003)	
	* PureCoverage 2003.06.00 Solaris 2 (32-bit), Copyright (C) 1994-2003 Rational Software Corp.	
	* All rights reserved.	
	* For contact information type: "purecov -help"	
	* Options <u>settings</u> : -purecov \ -purecov-home=/usr/rational/releases/purecov.sol.2003.06.00	
	* License successfully checked out	
Normal program output	* Command-line: a.out	
	Hello, World	
PureCoverage saves coverage data to a .pcv file	**** PureCoverage instrumented a.out (pid 3466) ****	
	* Saving coverage data to /usr/home/pat/example/a.out.pcv.	
	* To view results type: purecov -view /usr/home/pat/example/a.out.pcv	

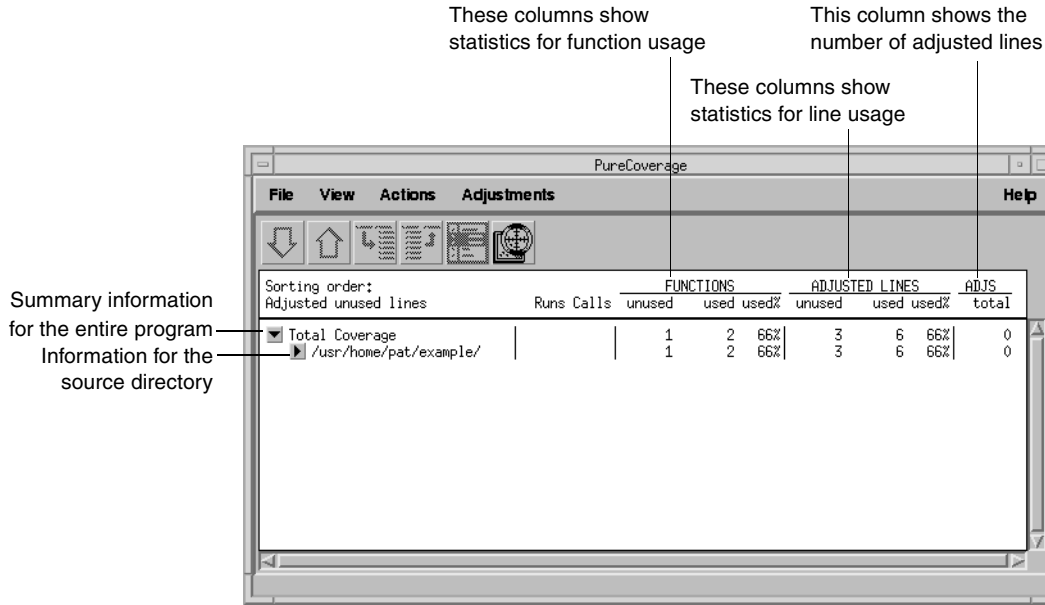
The `a.out` program produces its normal output, just as if it were not instrumented. When the program completes execution, PureCoverage writes coverage information for the session to the file `a.out.pcv`. Each time the program runs, PureCoverage updates this file with additional coverage data.

Displaying C/C++ coverage data

To display the coverage data for the program, use the command:

```
% purecov -view a.out.pcv
```

This displays the PureCoverage Viewer.

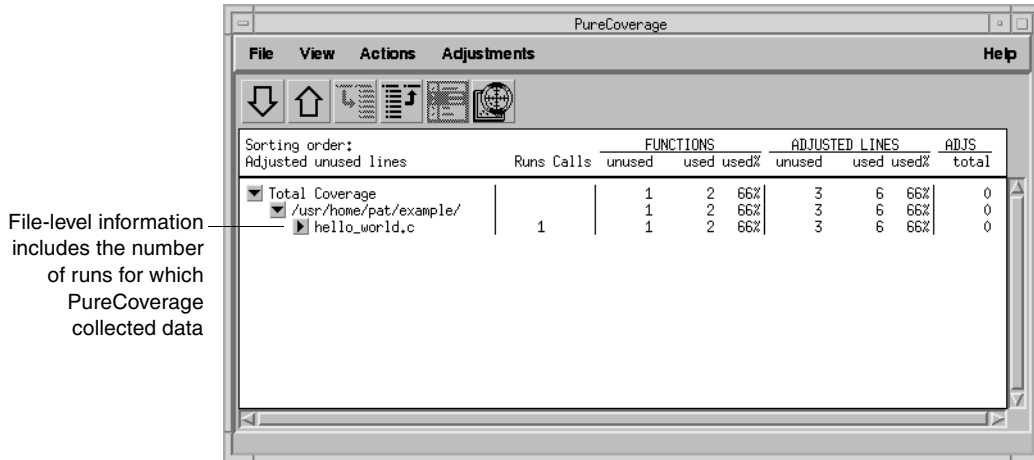


In this example, there is only one source directory, so the information displayed for the directory is identical to the `Total Coverage` information.

Note: The default header for line statistics is `ADJUSTED LINES`, not just `LINES`. This is because PureCoverage has an adjustment feature that lets you adjust coverage statistics by excluding specific lines. Under certain circumstances, the adjusted statistics give you a more practical reflection of coverage status than the actual coverage statistics. The `ADJS` column in this example contains zeroes, indicating that it does not include adjustments.

Expanding the file-level detail

Click ► next to `.../example/` to expand the file-level information for the directory.



You used only one file in the `example` directory to build `a.out`. Therefore the `FUNCTIONS` and `ADJUSTED LINES` information for the file is the same as for the directory. The number 1 in the `Runs` column indicates that you ran the instrumented `a.out` only once.

Note: When you are examining data collected for multiple executables, or for executables that have been rebuilt with some changed files, the number of runs can be different for each file.

Examining function-level detail

Expand the `hello_world.c` line to show function-level information.

The Viewer shows coverage information for the functions `display_message`, `main`, and `display_hello_world`.

The Calls column shows how many times the program called each function


The FUNCTIONS columns tell at a glance whether each function was used or unused

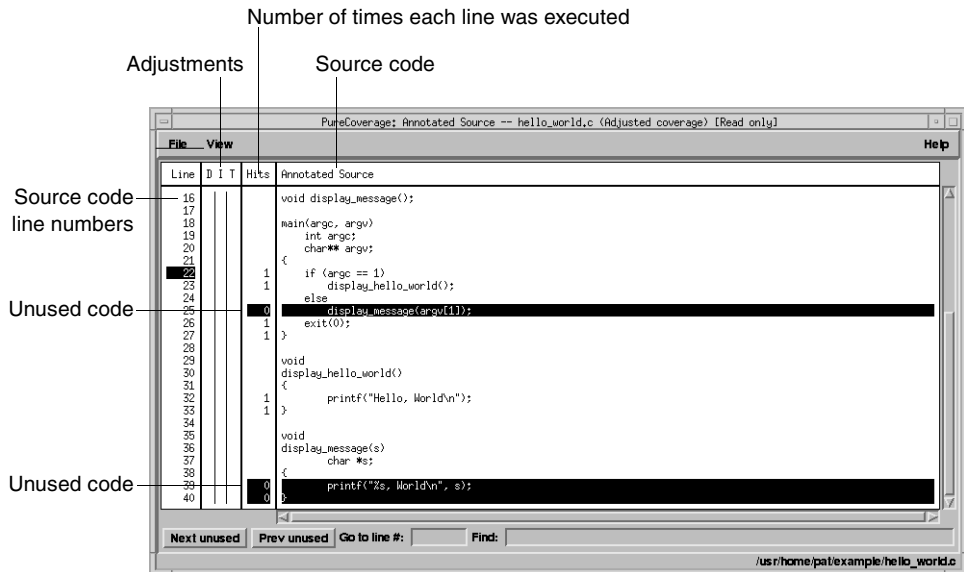
Function-level information includes the number of times the program called each function

Sorting order:		FUNCTIONS			ADJUSTED LINES			ADJS		
Adjusted	unused lines	Runs	Calls	unused	used	used%	unused	used	used%	total
▼	Total Coverage			1	2	66%	3	6	66%	0
▼	/usr/home/pat/example/			1	2	66%	3	6	66%	0
▼	hello_world.c	1		1	2	66%	3	6	66%	0
	display_message		0	unused			2	0	0%	0
	main		1		used		1	4	80%	0
	display_hello_world		1		used		0	2	100%	0

PureCoverage does not list the `printf` function or any functions that it calls. The `printf` function is a part of the system library, `libc`. By default, PureCoverage excludes collection of data from system libraries.

Examining the annotated source

To see the source code for `main` annotated with coverage information, click the Annotated Source tool  next to `main` in the Viewer. PureCoverage displays the Annotated Source window.



PureCoverage highlights code that was not used when you ran the program. In this file only two pieces of code were not used:

- The `display_message(argv[1]);` statement in `main`
- The entire `display_message` function

A quick analysis of the code reveals the reason: the program was invoked without arguments.

Improving Hello World's test coverage

To improve the test coverage for Hello World:

- 1 Without exiting PureCoverage, run the program again, this time with an argument. For example:

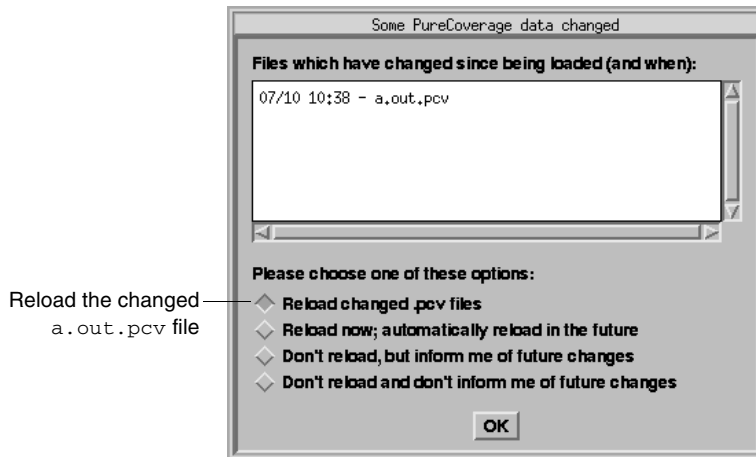
```
% a.out Goodbye
```

PureCoverage displays the following:

```
**** PureCoverage instrumented a.out (pid 17331 at Wed Feb
6 10:38:07 2003) PureCoverage 2003.06.00 Solaris 2
(32-bit), Copyright (C) 1994-2003 Rational Software Corp.
* All rights reserved.
* For contact information type: "purecov -help"
* Options settings: -purecov \
-purecov-home=/usr/rational/releases/purecov.sol.2003.
06.00
* License successfully checked out
* Command-line: a.out Goodbye
Goodbye, World
```

```
**** PureCoverage instrumented a.out (pid 17331) ****
* Saving coverage data to
/usr/home/pat/example/a.out.pcv.
* To view results type: purecov -view
/usr/home/pat/example/a.out.pcv
```

- 2 PureCoverage displays a dialog confirming that coverage data has changed for this run. Select **Reload changed .pcv files** and click **OK**.



Note: This dialog appears only if the PureCoverage Viewer is open when you run the program.

PureCoverage updates the coverage information in the Viewer and the Annotated Source window.

Function and line coverage is now 100%

The screenshot shows the PureCoverage application window with a menu bar (File, View, Actions, Adjustments, Help) and a toolbar. Below the toolbar is a table with the following data:

Adjusted unused lines	Runs	Calls	FUNCTIONS			ADJUSTED LINES			ADJS total
			unused	used	used%	unused	used	used%	
0			0	3	100%	0	9	100%	0
0			0	3	100%	0	9	100%	0
0	2		0	3	100%	0	9	100%	0
		1		used		0	2	100%	0
		1		used		0	2	100%	0
	2			used		0	5	100%	0

The statement
display_message
(argv[1]);...
and the function
display_message are
now shown as used

The screenshot shows the PureCoverage: Annotated Source window for the file hello_world.c. The window displays the following code with coverage hits indicated in the 'Hits' column:

Line	D	I	T	Hits	Annotated Source
16					void display_message();
17					
18					main(argc, argv)
19					int argc;
20					char** argv;
21					{
22				2	if (argc == 1)
23				1	display_hello_world();
24					else
25				1	display_message(argv[1]);
26				2	exit(0);
27				2	}
28					
29					void
30					display_hello_world()
31					{
32				1	printf("Hello, World\n");
33				1	}
34					
35					void
36					display_message(s)
37					char *s;
38					{
39				1	printf("%s, World\n", s);

Note: If you still have untested lines, it is possible that your compiler is generating unreachable code.

3 Select **File > Exit**.

Viewing UNIX coverage data on Windows

You can collect coverage data on your UNIX system and view it on Windows using Rational PureCoverage for Windows.

To collect coverage data for viewing on Windows, assign the value `windows` or `both` to the `-view-file-format` option. You can specify the option in the environment variable `PURECOVOPTIONS` or on the command line.

With the option set to `windows`, PureCoverage saves coverage data to a `.cfy` file, which you can analyze using PureCoverage for Windows. With the option set to `both`, PureCoverage saves data to a `.pcv` file as well.

PureCoverage for UNIX does not merge `.cfy` files, unlike `.pcv` files. You can merge `.cfy` files when you view them on Windows.

For more information, see the online help systems for PureCoverage on both UNIX and Windows.

Using report scripts

You can use PureCoverage report scripts to format and process PureCoverage data. The report scripts are located in the `<purecovhome>/scripts` directory.

Select **File > Run script** to open the script dialog.



You can also run report scripts from the command line.

Report scripts

`pc_annotate` Produces an annotated source text file

```
% pc_annotate [-force-merge] [-apply-adjustments=no] \  
[-file=<basename>...] [-type=<type>] [<prog>.pcv...]
```

`pc_below` Reports low coverage

```
% pc_below [-force-merge] [-apply-adjustments=no] [-percent=<pct>] \  
[<prog>.pcv...]
```

Report scripts

pc_build_diff Compares PureCoverage data from two builds of an application

```
% pc_build_diff [-apply-adjustments=no] [-prefix=XXXX... ] old.pcv \  
new.pcv
```

pc_covdiff Annotates the output of diff for modified source code

Note: Cannot run from Viewer

```
% yourdiff <name> | pc_covdiff [-context=<lines>] \  
[-format={diff|side-by-side|new-only}] [-lines=<boolean>] \  
[-tabs=<stops>] [-width=<width>] [-force-merge] [-apply-adjustments=no] \  
-file=<name> <prog>.pcv...
```

pc_diff Lists files for which coverage has changed

```
% pc_diff [-apply-adjustments=no] old.pcv new.pcv
```

pc_email Mails a report to the last person who modified insufficiently covered files

```
% pc_email [-force-merge] [-apply-adjustments=no] [-percent=<pct>] \  
[<prog>.pcv...]
```

pc_select Identifies the subset of tests required to exercise modified source code

```
% <list of changed files> | pc_select \  
[-diff=<rules>] [-canonicalize=<rule>] test1.pcv test2.pcv...
```

pc_ssheets Produces a summary in spreadsheet format

```
% pc_ssheets [-force-merge] [-apply-adjustments=no] [<prog>.pcv...]
```

pc_summary Produces an overall summary in table format

```
% pc_summary [-file=<name>...] [-force-merge] [-apply-adjustments=no] \  
[<prog>.pcv...]
```

PureCoverage options

PureCoverage provides command-line options for controlling operations and handling coverage data both for C/C++ and for Java code.

Build-time options

For a C or C++ application, specify build-time options on the link line when you instrument with PureCoverage. For example:

```
% purecov -always-use-cache-dir cc ...
```

For a Java application, specify these options (which for Java are not actually build-time options) on the command line when you run the application with PureCoverage.

For C, C++, and Java applications, you can also set these options using the `PURECOVOPTIONS` environment variable. For example:

```
% setenv PURECOVOPTIONS "-always-use-cache-dir"
```

	Commonly used build-time options	Default
*	<code>-always-use-cache-dir</code> Forces all PureCoverage instrumented object files to be written to the global cache directory	no
	<code>-auto-mount-prefix</code> Removes the prefix used by file system auto-mounters	/tmp_mnt
*	<code>-cache-dir</code> Specifies the global directory for caching instrumented object files	<purecovhome>/cache
*	<code>-collector</code> Specifies the collect program to handle static constructors (for use with gcc, g++)	none
*	<code>-ignore-runtime-environment</code> Prevents the runtime PureCoverage environment from overriding the option values used in building the program	no
*	<code>-linker</code> Specifies a linker other than the system default for building executables	system-dependent

* Does not apply to Java.

Runtime options

For a C or C++ application, specify runtime options on the link line when you instrument with PureCoverage. For a Java application, specify these options on the command line when you run the application with PureCoverage.

For C, C++, and Java applications, you can also set these options using the `PURECOVOPTIONS` environment variable. For example:

```
% setenv PURECOVOPTIONS \  
"-counts-file=./test1.pcv `printenv PURECOVOPTIONS`"
```

	Commonly used runtime options	Default
†	<code>-counts-file</code> Specifies an alternate file for writing coverage count data in binary format	<code>%v.pcv</code>
	<code>-follow-child-processes</code> Controls whether PureCoverage is enabled in forked child processes	<code>no</code>
†	<code>-log-file</code> Specifies a log file for PureCoverage runtime messages	<code>stderr</code>
*	<code>-program-name</code> Specifies the full pathname of the PureCoverage instrumented program	<code>argv[0]</code>
†	<code>-user-path</code> Specifies a list of directories to search for source code	<code>none</code>

† Can use the conversion characters listed on page 45.

* Does not apply to Java.

Analysis-time options

Use analysis-time options with analysis-time mode options. For example:

```
% purecov -merge=result.pcv -force-merge filea.pcv fileb.pcv
```

Commonly used analysis-time options	Default
-apply-adjustments Applies all adjustments in the <code>\$HOME/.purecov.adjust</code> file to exported coverage data	yes
-force-merge Forces the merging of coverage data files (<code>.pcv</code>) obtained from different versions of the same object file	no

Analysis-time mode options

Command-line syntax:

```
% purecov -<mode option> [analysis-time options] \  
<file1.pcv file2.pcv ...>
```

Analysis-time mode options	Compatible options
-export Merges and writes coverage counts from multiple coverage data files (<code>.pcv</code>) in export format to a specified file (<code>-export=<filename></code>) or to <code>stdout</code>	-apply-adjustments
-extract Extracts adjustment data from source code files and writes it to <code>\$HOME/.purecov.adjust</code>	none
-merge=<filename.pcv> Merges and writes coverage counts from multiple coverage data files (<code>.pcv</code>) in binary format	-force-merge
-view Opens the PureCoverage Viewer for analysis of one or more coverage data files (<code>.pcv</code>)	-force-merge, -user-path

Rational Quantify: What it does

Your application's runtime performance—its speed—is one of its most visible and critical characteristics. Developing high-performance software that meets the expectations of customers is not an easy task. Complex interactions between your code, third-party libraries, the operating system, hardware, networks, and other processes make identifying the causes of slow performance difficult.

Rational[®] Quantify[®] is a powerful tool that identifies the portions of your application that dominate its execution time. It supports C and C++ applications, as well as Java applications running on a Solaris SPARC 32-bit Java virtual machine (JVM). Quantify gives you the insight to eliminate performance problems so that your software runs faster. With Quantify, you can:

- Get accurate and reliable performance data
- Control how data is collected, collecting data for a small portion of your application's execution or the entire run
- Compare *before* and *after* runs to see the impact of your changes on performance
- Easily locate and fix only the problems with the highest potential for improving performance

This chapter introduces the basic concepts involved in using Quantify. For complete information, see the Quantify online help system, including the *Java Supplement* for Quantify.

How Quantify works: C/C++

Unlike sampling-based profilers, Quantify reports performance data for your program without any profiler overhead. The numbers you see represent the time your program would take without Quantify.

Quantify instruments and reports performance data for *all* the code in your program, including system and third-party libraries, shared libraries, and statically linked modules.

Quantify counts machine cycles: For C/C++ code, Quantify uses Object Code Insertion (OCI) technology to count the instructions your program executes and to compute how many cycles they require to execute. Counting cycles means that the time Quantify records in your code is independent of accidental local conditions and, assuming that the input does not change, identical from run to run. The fact that performance data is **repeatable** enables you to see precisely the effects of algorithm and data-structure changes.

Since Quantify counts cycles, it gives you accurate data at any scale. You do *not* need to create long runs or make numerous short runs to get meaningful data as you must with sampling-based profilers—one short run and you have the data. As soon as you can run a test program, you can collect meaningful performance data and establish a baseline for future comparison.

Quantify times system calls: Quantify measures the elapsed (wall clock) time of each system call made by your program and reports how long your program waited for those calls to complete. You can immediately see the effects of improved file access or reduced network delay on your program. You can optionally choose to measure system calls by the amount of time the kernel records for the process, which is the same as the time the UNIX `/bin/time` utility records.

Quantify distributes time accurately: Quantify distributes each function's time to its callers so you can tell at a glance which function calls were responsible for the majority of your program's time. Unlike `gprof`, Quantify does not make assumptions about the average cost per function. Quantify measures it directly.

How Quantify works: Java

Quantify provides performance profiling for Java applications running on a Solaris SPARC 32-bit Java virtual machine (JVM).

Quantify times performance: Quantify times each method as it executes, and collects accurate data about the actual execution of your Java code. You can choose either to record elapsed wall-clock time or to measure the amount of time the kernel records for the process, like the

UNIX `/bin/time` utility. Because data for Java code is based on timing and not counting cycles, as it is for C and C++, performance data for Java code, while reliable for a given run, is not repeatable.

Quantify distributes time accurately: Quantify distributes each method's time to its callers. This helps you detect the methods that are ultimately responsible for bottlenecks in your code.

Collecting performance data: C/C++

To collect performance data for a C/C++ program:

1 Add `quantify` to the front of the *link* command line. For example:

```
% quantify cc -g hello_world.c -o hello_world
```

2 Run the instrumented program as you usually do:

```
% hello_world
```

When the program starts, Quantify prints license and support information, followed by the expected output from your program.

```
**** Quantify instrumented hello_world (pid 20352 at Sat 5
08:41:27 1999)
Quantify 4.4 Solaris 2, Copyright 1993-1999 Rational
Software Corp.
* For contact information type: "quantify -help"
* Quantify licensed to Quantify Evaluation User
* Quantify instruction counting enabled.
```

Program output—Hello, World.

Data transmission—Quantify: Sending data for 37 of 1324 functions
from hello_world (pid 20352).....done.

When the program finishes execution, Quantify transmits the performance data it collected to `qv`, Quantify's data-analysis program.

Interpreting the program summary: C/C++

After each dataset is transmitted, Quantify prints a program summary showing at a glance how the original, non-instrumented, program is expected to perform.

Time Quantify expects the original program to take

```
Quantify: Resource Statistics for hello_world (pid 20352)
*
*      cycles      secs
* Total counted time:      16148821      0.323 (100.0%)
*   Time in your code:           2721      0.000 (  0.0%)
*   Time in system calls:      843950      0.017 (  5.2%)
*   Dynamic library loading:  15302150      0.306 ( 94.8%)
*
* Note: Data collected assuming a sparcstation_lx with clock rate of 750 MHz.
* Note: These times exclude Quantify overhead and possible memory effects.
*
* Elapsed data collection time:      0.336 secs
*
* Note: This measurement includes Quantify overhead.
```

Time spent executing program functions (compute-bound) ————

Time spent waiting for system calls to complete ————

Time spent loading dynamic libraries ————

Time taken to collect data includes Quantify's counting overhead and any memory effects ————

Collecting performance data: Java

To collect Java performance data, run Quantify with the `-java` option, as follows:

- For an applet:

```
% quantify [<Quantify options>] -java <applet viewer>
[<applet viewer options>] <html file>
```

- For a class file:

```
% quantify [<Quantify options>] -java <Java executable>
[<Java options>] <class>
```

- For a JAR file:

```
% quantify [<Quantify options>] -java <Java executable>
[<Java options>] -jar <JAR file>
```

- For a container program such as Netscape Navigator:

```
% quantify [<Quantify options>] -java <exename>
[<arguments to exename>]
```

Note: Quantify can collect line-by-line performance data or method-level data. By default, Quantify uses the line level when debug data, which is stored in class files, is available.

When Quantify starts, it prints license and support information, followed by the expected output from your program.

When the program finishes execution, Quantify transmits the performance data it collected to `qv`, Quantify's data-analysis program.

Interpreting the program summary: Java

After each dataset is transmitted, Quantify prints a program summary showing at a glance how the original, non-instrumented, program is expected to perform.

Actual time taken by your process
from when Java was started to the end

Time spent in your Java code, including some Quantify overhead

```
Quantify: Resource Statistics for /people/jo/4java/bin/appletviewer (pid 24565)
*
*           cycles           secs
* Total counted time:      101071710000    134.762 (100.0%)
* Time in your code:       54225003750     72.300 ( 53.7%)
*
*Time Quantify excluded from the dataset:
* Time spent blocked/waiting: 36447111000    48.596
*
* Note: Data collected assuming a UltraSparc with clock rate of 750 MHz.
*
* Note: These measurements performed on a machine with 2 processors.
* For threaded programs on multiprocessors, Quantify will time
* operations that are executed in parallel as if they were performed
* on a single processor.
*
```

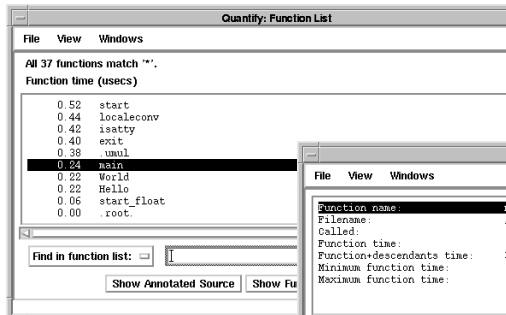
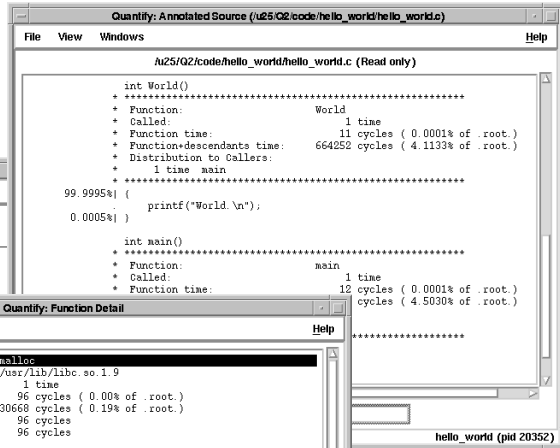
Using Quantify's data analysis windows

After transmitting the last dataset, Quantify displays the Control Panel. From here, you can display Quantify's data analysis windows and begin analyzing your program's performance.

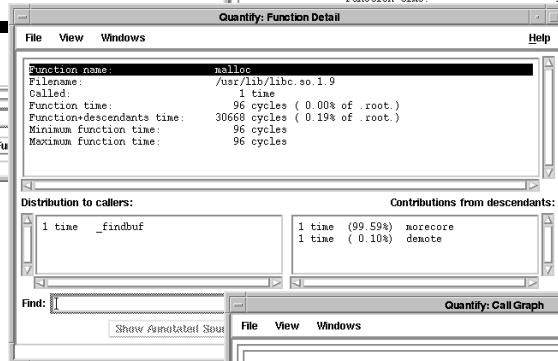
CONTROL PANEL



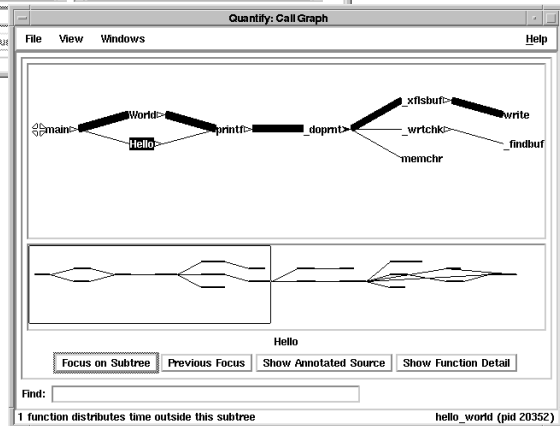
ANNOTATED SOURCE



FUNCTION LIST



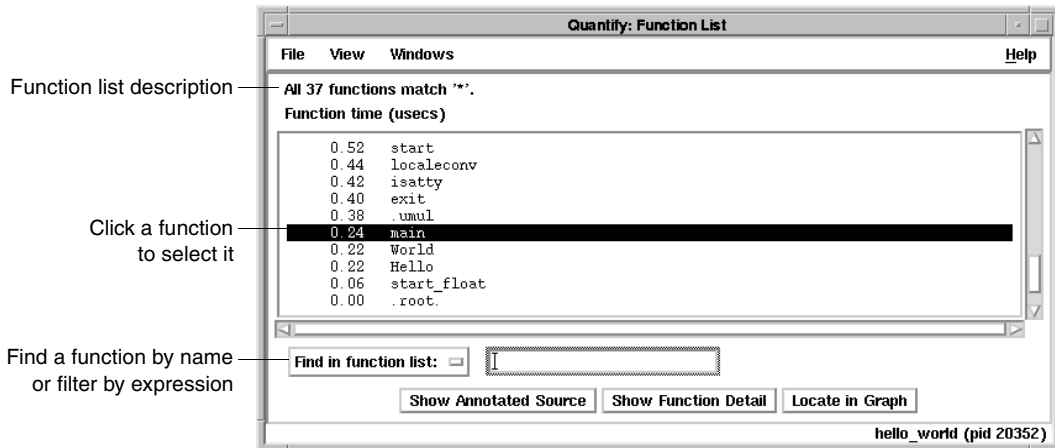
FUNCTION DETAIL



CALL GRAPH



The Function List window

The Function List window shows the functions that your program executed. By default, it displays all the functions in your program, sorted by their *function time*. *This is the amount of time a function spent performing computations (compute-bound) or waiting for system calls to complete.*



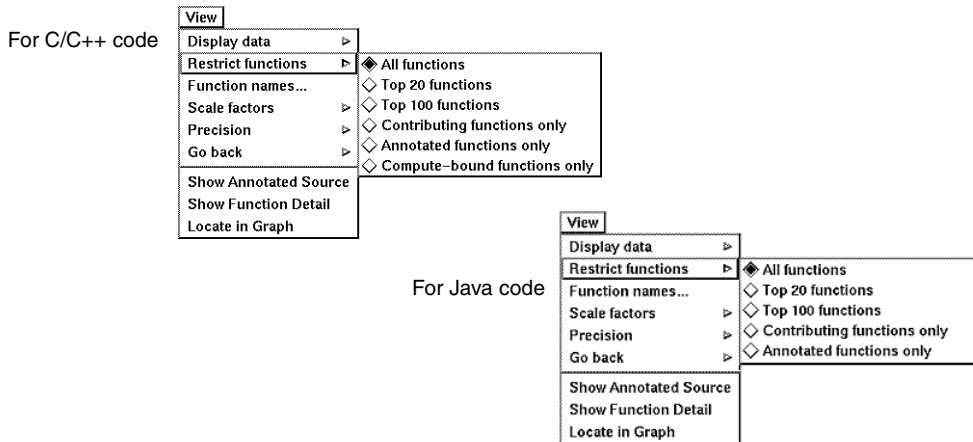
Sorting the function list

To sort the function list based on the various data Quantify collects, select **View > Display data**.



Restricting functions

To focus attention on specific types of functions, or to speed up the preparation of the function list report in large programs, you can restrict the functions shown in the report. Select **View > Restrict functions**.

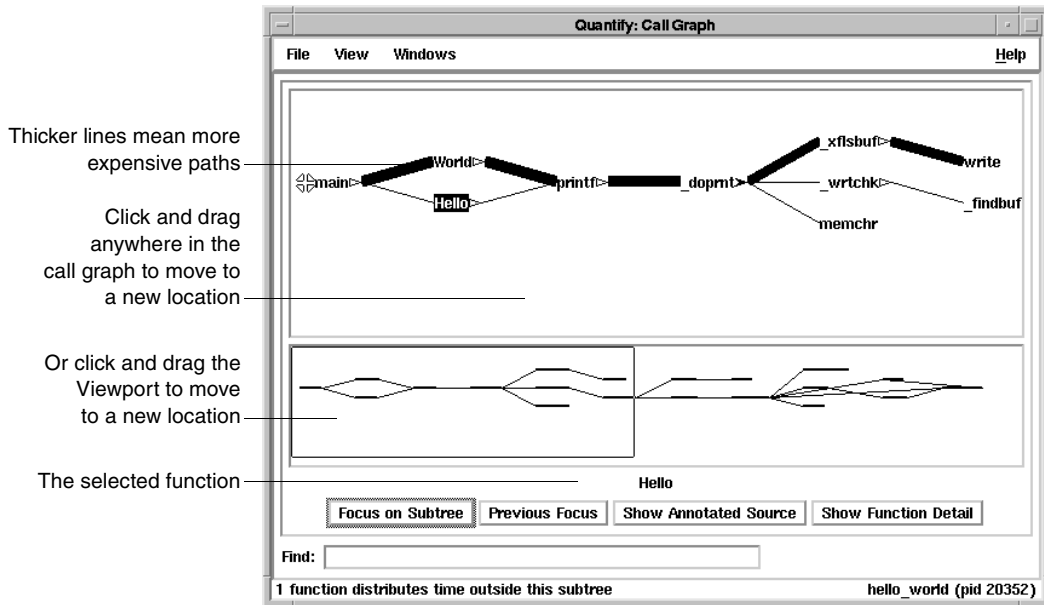


You can restrict the list to the top 20 or top 100 functions in the list, to the functions that have annotated source, to functions that are compute-bound (make no system calls), or to functions that contribute non-zero time for a recorded data type.

The Call Graph window

The Call Graph window presents a graph of the functions called during the run. It uses lines of varying thickness to graphically depict where your program spends its time. Thicker lines correspond directly to larger amounts of time spent along a path.


The call graph helps you understand the calling structure of your program and the major call paths that contributed to the total time of the run. Using the call graph, you can quickly discover the sources of bottlenecks.



By default, Quantify expands the call paths to the top 20 functions contributing to the overall time of the program.

Using the pop-up menu

To display the pop-up menu, right-click any function in the call graph.




You can use the pop-up menu to:

- Expand and collapse the function's subtree
- Locate individual caller and descendant functions
- Change the focus of the call graph to the selected function
- Display the annotated source code or the function detail for the selected function

Expanding and collapsing descendants

Use the pop-up menu to expand or collapse the subtrees of descendants for individual functions.



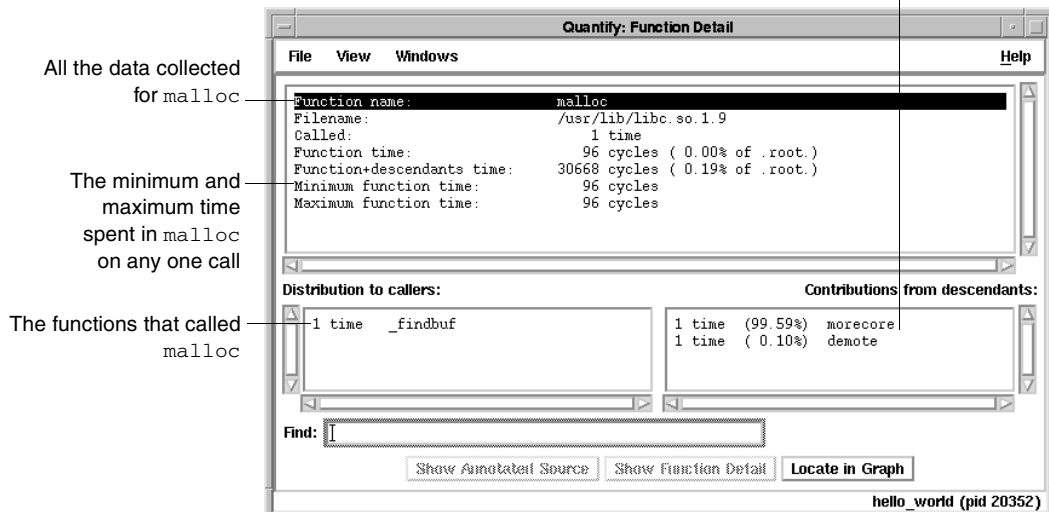
After expanding or collapsing subtrees, you can select **View > Redo layout** to remove any gaps that your changes create in the call graph.

The Function Detail window

The Function Detail window presents detailed performance data for a single function, showing its contribution to the overall execution of the program.

For each function, Quantify reports both the time spent in the function's own code (its *function* time) and the time spent in all the functions that it called (its *descendants* time). Quantify distributes this accumulated *function+descendants* time to the function's immediate caller.

The immediate descendants of `malloc`, and how they contributed to `malloc`'s function+descendants time




Double-click a caller or descendant function to display the detail for that function.

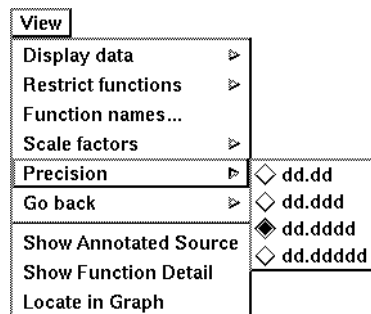
The function time and the function+descendants time are shown as a percentage of the total accumulated time for the entire run. These percentages help you understand how this function's computation contributed to the overall time of the run. These times correspond to the thickness of the lines in the call graph.

Changing the scale and precision of data

Quantify can display the recorded data in cycles (the number of machine cycles) and in microseconds, milliseconds, or seconds. To change the scale of data, select **View > Scale factors**.



To change the precision of data, select **View > Precision**.



Saving function detail data

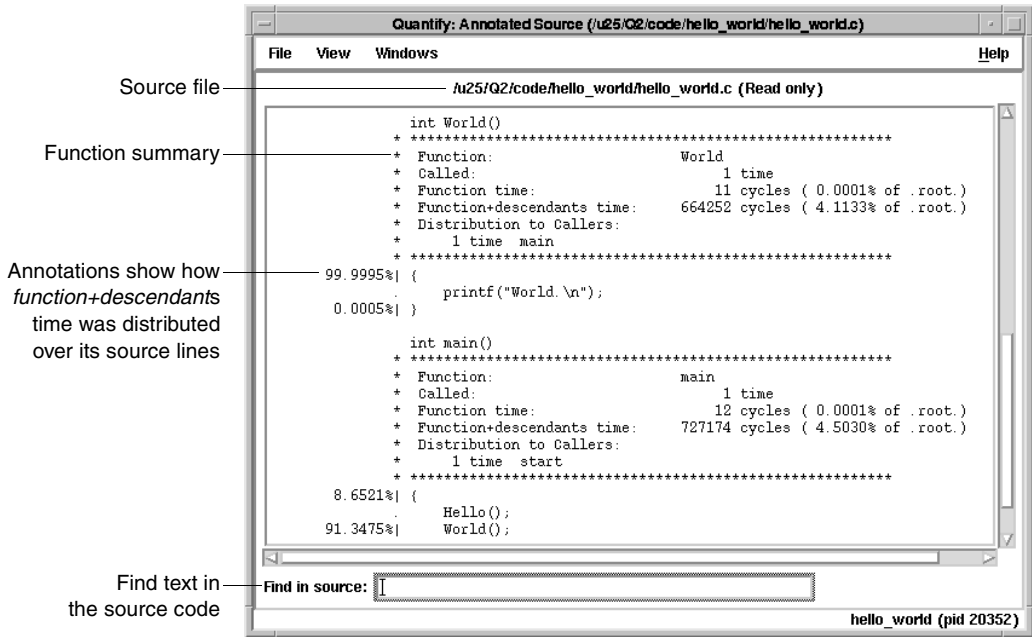
To save the current function detail display to a file, select **File > Save current function detail as**.

To append additional function detail displays to the same file, select **File > Append to current detail file**.

The Annotated Source window

Quantify's Annotated Source window presents line-by-line performance data using the function's source code.


Note: The Annotated Source window is available only for files that you compile using the `-g` debugging option.



The numeric annotations in the margin reflect the time recorded for that line or basic block over all calls to the function. By default, Quantify shows the function time for each line, scaled as a percentage of the total function time accumulated by the function.


Changing annotations

To change annotations, use the View menu. You can select both *function* and *function+descendants* data, either in cycles or seconds and as a percentage of the *function+descendants* time.



Saving performance data on exit

To exit Quantify, select **File > Exit Quantify**. If you analyze a dataset interactively, Quantify does not automatically save the last dataset it receives. When you exit, you can save the dataset for future analysis.



By default, Quantify names dataset files to reflect the program name and its runtime process identifier. You can analyze a saved dataset at a later time by running `qxv`, Quantify's data analysis program.

You can also save Quantify data in export format. This is a clear-text version of the data suitable for processing by scripts.

Comparing program runs with `qxdiff`

The `qxdiff` script compares two export data files and reports any changes in performance. For C or C++ programs, the results show exactly how much your program's performance has improved. For Java code, the results indicate general performance trends. This is because C and C++ performance data, based on counting cycles, is repeatable, while Java data, based on the timing of methods, is not repeatable.

To use the `qxdiff` script:

- 1 Save baseline performance data to an export file. Select **File > Export Data As** in any data analysis window.
- 2 Change the program and run Quantify on it again.
- 3 Select **File > Export Data As** to export the performance data for the new run.
- 4 Use the `qxdiff` script to compare the two export data files. For example:

```
% qxdiff -i testHash.pure.20790.0.qx  
improved_testHash.pure.20854.0.qx
```

You can use the `-i` option to ignore functions that make calls to system calls.

Below is the output from this example.

```
Differences between:
program testHash.pure (pid 20790) and
program improved_testHash.pure (pid 20854)

qxdiff lists the
functions that have
changed ...
and summarizes the
differences for the
entire run
```

Function name	Calls	Cycles	% change
strcmp	-40822	-1198640	93.77% faster
putHash	0	-32912	6.61% faster
getHash	0	-28376	7.86% faster
remHash	0	-7856	5.91% faster
hashIndex	0	10000	1.49% slower

5 differences; -1257784 cycles (-0.025 secs at 50 MHz)
25.01% faster overall (ignoring system calls).

Quantify options

Quantify provides command-line options for controlling operations and handling coverage data for both C/C++ and for Java code.

Build-time options

For a C or C++ application, specify build-time options on the link line when you instrument with Quantify. For example:

```
% quantify -cache-dir=$HOME/cache -always-use-cache-dir \
cc ...
```

For a Java application, specify these options (which for Java are not actually build-time options) on the command line when you run the application with Quantify.

For C, C++, and Java applications, you can also set these options using the `QUANTIFYOPTIONS` environment variable. For example:

```
% setenv QUANTIFYOPTIONS "-always-use-cache-dir"
```

Commonly used build-time options	Default
-always-use-cache-dir Specifies whether instrumented files are written to the global cache directory	no
-cache-dir Specifies the global cache directory	<quantifyhome>/cache

	Commonly used build-time options	Default
	-collection-granularity Specifies the level of collection granularity	line
†	-collector Specifies the collect program to handle static constructors in C++ code	none
†	-ignore-runtime-environment Prevents the runtime Quantify environment from overriding option values used in building the program	no
†	-linker Specifies an alternative linker to use instead of the system linker	system-dependent
†	-use-machine Specifies the build-time analysis of instruction times according to a particular machine	system-dependent

† Does not apply to Java.

qv runtime options

To run `qv`, specify the option and the saved `.qv` file. For example:

```
% qv -write-summary-file a.out.23.qv
```

	qv options	Default
	-add-annotation Specifies a string to add to the binary file	none
	-print-annotations Writes the annotations to <code>stdout</code>	no
	-windows Controls whether Quantify runs with the graphical interface	yes
	-write-export-file Writes the recorded data in the dataset to a file in export format	none
	-write-summary-file Writes the program summary for the dataset to a file	none

Runtime options

For a C or C++ application, specify build-time options on the link line when you instrument with Quantify. For a Java application, specify these options on the command line when you run the application with Quantify.

For C, C++, and Java applications, you can also set these options using the `QUANTIFYOPTIONS` environment variable. For example:

```
% setenv QUANTIFYOPTIONS "--windows=no"; a.out
```

	Commonly used runtime options	Default
†	-avoid-recording-system-calls Avoids recording specified system calls	system-dependent
	-measure-timed-calls Specifies measurement for timing system calls	elapsed-time
†	-record-child-process-data Records data for child processes created by <code>fork</code> and <code>vfork</code>	no
†	-record-system-calls Records system calls	yes
†	-report-excluded-time Reports time that was excluded from the dataset	0.5
	-run-at-exit Specifies a shell script to run when the program exits	none
	-run-at-save Specifies a shell script to run each time the program saves counts	none
	-save-data-on-signals Saves data on fatal signals	yes
	-save-thread-data Saves composite or per-stack thread data	composite
	-write-export-file Writes the dataset to an export file as ASCII text	none
	-write-summary-file Writes the program summary for the dataset to a file	/dev/tty

Commonly used runtime options	Default
-windows Specifies whether Quantify runs with the graphical interface	yes

† Does not apply to Java.

API functions: C/C++

To use Quantify API functions with C/C++ code, include `<quantifyhome>/quantify.h` in your code and link with `<quantifyhome>/quantify_stubs.a`

Commonly used C/C++ functions	Description
<code>quantify_help (void)</code>	Prints description of Quantify API functions
<code>quantify_is_running (void)</code>	Returns <code>true</code> if the executable is instrumented
<code>quantify_print_recording_state (void)</code>	Prints the recording state of the process
<code>quantify_save_data (void)</code>	Saves data from the start of the program or since last call to <code>quantify_clear_data</code>
<code>quantify_save_data_to_file (char * filename)</code>	Saves data to a file you specify
<code>quantify_add_annotation (char * annotation)</code>	Adds the specified string to the next saved dataset
<code>quantify_clear_data (void)</code>	Clears the performance data recorded to this point
† <code>quantify_<action>_recording_data (void)</code>	Starts and stops recording of all data
† <code>quantify_<action>_recording_dynamic_library_data (void)</code>	Starts and stops recording dynamic library data
† <code>quantify_<action>_recording_register_window_traps (void)</code>	Starts and stops recording register-window-trap data
† <code>quantify_<action>_recording_system_call (char *system_call_string)</code>	Starts and stops recording specific system-call data

Commonly used C/C++ functions	Description
† <code>quantify_<action>_recording_system_calls</code> (void)	Starts and stops recording of all system-call data

† *<action>* is one of: start, stop, is. For example:
`quantify_stop_recording_system_call`

API methods: Java

You can call an API method from your Java code or from a debugger. Use the following syntax:

```
Rational.PureAPI.IsRunning()
```

or

```
import Rational.PureAPI;
...
PureAPI.IsRunning()
```

`PureAPI` is a Java class that includes all the Quantify API methods that can be used with Java code. The `PureAPI` class is part of a Java package called `Rational.jar`, which is located in `<quantifyhome>`.

You can run class files that include calls to `PureAPI` methods with or without Quantify:

- When you run these class files with Quantify, Quantify automatically sets `CLASSPATH` and `LD_LIBRARY_PATH` to access `Rational.jar` and `libQProfJ.so`.
- When you run the class files without Quantify, you must add `<quantifyhome>/lib32` to your `LD_LIBRARY_PATH`. In addition, if you do not have a `Rational.jar` file in your `<javahome>/jre/lib/ext` directory, you must add `<quantifyhome>` to your `CLASSPATH`.

The Java API methods are as follows:

Java API methods: class <code>PureAPI</code>	Description
<code>public static int IsRunning();</code>	Returns <code>true</code> if the executable is instrumented
<code>public static int DisableRecordingData();</code>	Disables collection of all data by Quantify

Java API methods: class PureAPI	Description
<code>public static int StartRecordingData();</code>	Tells Quantify to start recording all program performance data
<code>public static int StopRecordingData();</code>	Tells Quantify to stop recording all program performance data
<code>public static int IsRecordingData();</code>	Checks if Quantify is currently recording all program performance data
<code>public static int ClearData();</code>	Tells Quantify to clear all the data it has recorded about your program's performance to this point
<code>public static int SaveData();</code>	Saves all the data recorded since program start (or the last call to <code>clearData()</code>) into a dataset (a <code>.qv</code> file)
<code>public static int AddAnnotation(String annotation);</code>	Tells Quantify to save the argument string in the next output datafile written by <code>saveData()</code>

Index

Symbols

%V, %v, %p 47

A

ABR, array bounds read error
 correcting 38
 in Hello World 36
access errors, how Purify finds 50
account number, Rational Software 4
AccountLink user input 4
-add-annotation 83
adjusted lines 58
-always-use-cache-dir 46, 66, 82
analysis-time options 67
Annotated Source window
 PureCoverage 61
 Quantify 80
a.out.pcv 57
API functions
 Purify 45
 Quantify (C/C++) 85
 Quantify (Java) 86
appending function detail 79
applets, collecting performance data 72
-apply-adjustments 68
-auto-mount-prefix 47, 66
-avoid-recording-system-calls 84

B

blue memory color 51
build-time options
 PureCoverage 66
 Purify 46
 Quantify 82

C

cache directory
 configuring 19
 location of 25
-cache-dir 19, 46, 66, 82, 83
caching dynamic shared objects
 IRIX 33
caching options
 PureCoverage 66
 Purify 46
 Quantify 82
Call Graph window, Quantify 78
Calls column, PureCoverage 60
CD-ROMs
 ejecting 26
 mounting 25
.cfy Windows coverage data file 64
-chain-length 47
changing annotations, Quantify 80
characters, conversion 47
class files, collecting performance data 72
ClearQuest, integrating 18
code, see source code
collapsing subtrees 78
-collection-granularity 83
-collector 46, 66, 83
color, see memory color
comparing program runs
 with PureCoverage 61
 with Purify 42
 with Quantify qxdiff script 81
compiling and linking 33
compute-bound
 functions 75, 76
 time 72
configuration message 35
configure command 27
configuring the cache directory 18

- container programs
 - collecting Java coverage data 54
 - collecting Java performance data 73
- controls, Purify program 35
- conversion characters for filenames 47
- counts-file 67
- coverage data
 - file level 59
 - function level 59
 - in PureCoverage Viewer 57
- cycles
 - counted by Quantify 70
 - scale factor 79

D

- daemons, and licensing 28
- .dat license file 7, 10, 28
- data
 - comparing export files 81
 - saving Quantify data 81
- debugger(s)
 - JIT debugging 44
 - scripts on HP-UX 21
 - using with Purify 44
- debugging option, see -g debugging option
- deleting product releases 23
- directories
 - cache 19
 - installation 6, 9, 12, 13, 24
 - PureLA 7, 10
 - Rational 24–25
- disk space requirements 5
- dynamic library, timing 72
- dynamic shared object (DSO) caching 33

E

- editing source code 38, 40
- ejecting CD-ROMs 26
- environment variables
 - LM_LICENSE_FILE 28
 - MANPATH 19

- PATH 19
- PURECOVOPTIONS 66, 67
- PUREOPTIONS 19
- PURIFYOPTIONS 47
- QUANTIFYOPTIONS 82, 84
- executables, instrumenting (IRIX) 33
- expanding subtrees 78
- expiration date, licenses 13
- export 68
- exporting Quantify data 81
- extract 68

F

- fds-in-use-at-exit 47
- file(s)
 - a.out.pcv 57
 - .cfy Windows coverage data file 64
 - installing product 25
 - license_for_*.upd 5
 - Purify view 43
 - qv and qx script 20
 - Rational license (.dat) 7, 10, 28
 - rational.opt 22
 - Temporary.dat 28
 - users.purela 7, 10
- filename conversion characters 47
- filesystems, installing on read-only 18
- FLEXlm
 - commands 29
 - End User Manual 29
 - GLOBEtrouter web site 29
 - License Manager 27
- floating license 4
- follow-child-processes 48, 67
- force-merge 68
- Function Detail window 78
 - saving data 79
 - scale and precision of data 79
- Function List window
 - finding top contributors 75
 - restricting functions 76
- function+descendants time 78

- functions
 - compute-bound 76
 - coverage detail 59
 - restricting display in Quantify 76
 - sorting in Quantify 75
 - See also API functions
- Functions columns, PureCoverage 60

G

- g debugging option
 - and PureCoverage 56
 - and Purify 36
 - and Quantify 71, 80
- get_hostinfo.sh 4
- GLOBEtrouter web site 29
- graph, see Call Graph window
- green memory color 51

H

- heap analysis, Purify 41
- Hello World example
 - PureCoverage 55
 - Purify 32
- help system, setting up 18
- help, technical support x
- hiding
 - functions in Quantify 76
 - messages in Purify 43
- hostid for license server host 4
- HP-UX debugger scripts 21

I

- ignore-runtime-environment 46, 66, 83
- installation
 - basic steps 1
 - directory 6, 9, 12, 13, 24
 - evaluation license 3
 - on read-only filesystems 18
 - permanent license 24
 - requirements 5

- rs_install commands 27
- startup license 17
- Typical installation option 1
- user input (AccountLink) 4
- instrumenting a program
 - description of ix
 - with PureCoverage 56
 - with Purify 33
 - with Quantify 71
- integration, Purify and PureCoverage 45
- IRIX
 - compile/link command 33
 - DSO caching 33
 - running a Purify instrumented program 34

J

- JAR files, collecting performance data 72
- java
 - PureCoverage 54
 - Quantify 72
- Java and PureCoverage 54
- Java and Quantify 70, 72
- jit-debug 48
- just-in-time debugging 44

K

- keys, product license 17

L

- leaks, see memory leaks
- leaks-at-exit 48
- library
 - system and PureCoverage 60
 - time loading dynamic 72
- license daemon, lmgrd 28
- license file
 - .dat 7, 10, 22, 28
 - .upd 5
- license key types 17
- License Manager, FLEXlm 27

- license server
 - port number 9
 - requirements 4
 - using multiple servers 17
- license(s)
 - checking 27
 - expiration date 13
 - floating 4
 - license key types 17
 - named user 7, 10, 17, 19
 - permanent 24
 - quantity 4
 - Rational license file (.dat) 28
 - setting up 27
 - startup 2, 17
 - upd license file 5
 - user names 7, 10, 22–23
- license_check command 27
- license_for_*.upd file 5
- license_setup command 27
- line numbers
 - g option 33, 36
 - on IRIX 36
- linker 47, 66, 83
- links, symbolic 19
- LM_LICENSE_FILE environment variable 28
- lmgrd license daemon 28
- local variable names, displaying 33
- log-file 48, 67

M

- machine cycles 70
- MANPATH environment variable 19
- manual pages 19
- measure-timed-calls 84
- memory access errors
 - example 36
 - how Purify finds 50
- memory color 50
- memory in use message 41
- memory leaks
 - definition 41
 - heap analysis 41

- message 39
- new leaks button 39
- potential 41
- purify_new_leaks 45
- menu, Quantify pop-up 77
- merge 68
- messages 48
- messages
 - Purify 49
 - suppressing Purify 43
- MLK, memory leak 40
 - example 39
- mounting CD-ROMs 25

N

- Named User license 7, 10, 17, 19, 22
- new memory leaks, Purify 39

O

- Object Code Insertion (OCI) 70
- online help system, setting up 18
- operating system, identifying 25
- options
 - PureCoverage analysis-time 67
 - PureCoverage build-time 66
 - PureCoverage runtime 67
 - Purify build-time 46
 - Purify runtime 47
 - Quantify build-time 82
 - Quantify runtime 84
 - qv runtime 83
- options (by name)
 - add-annotation 83
 - always-use-cache-dir 46, 66, 82
 - apply-adjustments 68
 - auto-mount-prefix 47, 66
 - avoid-recording-system-calls 84
 - cache-dir 46, 66, 82, 83
 - chain-length 47
 - collection-granularity 83
 - collector 46, 66, 83
 - counts-file 67

- export 68
- extract 68
- fds-in-use-at-exit 47
- follow-child-processes 48, 67
- force-merge 68
- ignore-run-time-environment 66
- ignore-runtime-environment 46, 83
- java 54, 72
- jit-debug 48
- leaks-at-exit 48
- linker 47, 66, 83
- log-file 48, 67
- measure-timed-calls 84
- merge 68
- messages 48
- print-annotations 83
- print-home-dir 47
- program-name 48, 67
- record-child-process-data 84
- record-system-calls 84
- report-excluded-time 84
- run-at-exit 84
- run-at-save 84
- save-data-on-signals 84
- save-thread-data 84
- show-directory 48
- show-pc 48
- show-pc-offset 48
- use-machine 83
- user-path 48, 67
- view 54, 57, 68
- view-file 48
- windows 48, 83, 85
- write-export-file 83, 84
- write-summary-file 83, 84
- options file 22
- options_setup command 27
- overhead, Quantify 72

P

- PATH environment variable 19
- performance data 71, 73
 - saving 81

- permanent licenses
 - installing manually 24
 - requesting 24
- pop-up menu, Quantify 77
- port number, license server 9
- post_install command 27
- post-installation configuration tasks 18
- potential memory leak 41
- print-annotations 83
- print-home-dir 47
- product license keys 17
- producthome directory 24
- products, removing 23
- program controls, Purify 35
- program runs, comparing
 - Quantify qxdiff script 81
 - with PureCoverage 61
 - with Purify 42
- program summary, Quantify 72, 73
- program-name 48, 67
- programs, running instrumented
 - PureCoverage 56
 - Purify 34
 - Quantify 71
- purecov.configure command 27
- PureCoverage
 - benefits 53
 - for Windows 64
 - symbolic links for 20
 - using with Purify 45
 - Viewer 57
 - with Java 54
- PURECOVOPTIONS environment variable 67
- PureLA directory 7, 10
- PUREOPTIONS environment variable 19
- Purify
 - API functions 45
 - instrumenting a program 33
 - messages 49
 - Viewer 34
- purify.configure command 27
- PURIFYOPTIONS environment variable 47
- PurifyPlus ix, 1

Q

Quantify

- API functions (C/C++) 85
 - API functions (Java) 86
 - build-time options 82
 - Call Graph window 76, 78
 - overhead 72
 - repeatability of timing 70
 - runtime options 84
 - symbolic links for 20
 - with Java 70, 72
- quantify.configure command 27
- QUANTIFYOPTIONS environment variable 66, 82, 84
- qv 71, 73
- qv script files 20
- qx script files 20
- qxdiff script 81

R

- Rational ClearQuest, integrating 18
- rational daemon 28
- Rational license file
- *.dat 28
 - license_for_*.upd 5
- Rational PureCoverage for Windows 64
- Rational PurifyPlus ix, 1
- Rational Software account number 4
- Rational Software website
- AccountLink 5
 - home page x
 - technical support x
- rational.opt options file 22
- README file location 15
- read-only filesystems 18
- record-child-process-data 84
- record-system-calls 84
- red memory color 51
- Redo layout, Quantify 78
- removing previous releases 23
- report(s)
- program summary 72, 73
 - PureCoverage scripts 64

- report-excluded-time 84
- restricting functions in Quantify 76
- rs_install
 - commands 27
 - program 15
- run-at-exit 84
- run-at-save 84
- running an instrumented program
 - PureCoverage 56
 - Purify 34
 - Quantify 71
- runs
 - column, PureCoverage 59
 - comparing with PureCoverage 61
 - comparing with Purify 42
 - comparing with Quantify 81
- runtime options
 - PureCoverage 67
 - Purify 47
 - Quantify 84
 - qv 83

S

- save-data-on-signals 84
- save-thread-data 84
- saving
 - function detail data 79
 - Purify run 43
 - Quantify data 81
- scale factors 79
- scripts
 - enabling PureCoverage scripts 20
 - HP-UX debugger 21
 - PureCoverage report scripts 64
 - Quantify 20
 - qxdiff 81
- server-name.dat file 28
- servers, license 16
 - requirements 4
 - using multiple 17
- show-directory 48
- show-pc 48
- show-pc-offset 48

- sorting function list 75
- source code
 - annotated in PureCoverage 61
 - annotated in Quantify 80
 - displaying filenames 36
 - editing from Viewer 38, 40
 - line numbers, Purify 36
 - number of lines displayed 38
- startup license 2, 17
- statically allocated memory 51
- subtrees, Quantify 78
- summary, Quantify program 72, 73
- support, technical x
- suppressing Purify messages 43
- symbolic links 19
 - for HP-UX debugger scripts 21
 - for PureCoverage 20
 - for Purify 20
 - for Quantify 20
- system call timing 70
- system libraries and PureCoverage 60

T

- technical support x
- Temporary .dat file 28
- time
 - compute-bound 72
 - function+descendants 78
 - in code 72
 - loading dynamic libraries 72
 - to collect the data 72
- Total Coverage row, PureCoverage 58
- Typical installation option 1

U

- uname command 25
- uninstall command 23
- .upd license file 5
- use-machine 83

- user names, for licensing 7, 10, 22–23
- user-path 48, 67

V

- validating setup 18
- variable, see environment variable
- view 54, 57, 68
- view file, Purify 43, 44
- Viewer 57
 - PureCoverage 57
 - Purify 34
- view-file 48
- viewport, call graph 77

W

- websites
 - for obtaining Rational licenses 5
 - GLOBEtrotter 29
 - Rational software x
 - Rational technical support x
- windows 48, 83, 85
- windows
 - PureCoverage annotated source 61
 - PureCoverage viewer 57
 - Purify viewer 34
 - Quantify annotated source 80
 - Quantify call graph 78
 - Quantify data analysis 74
 - Quantify function detail 78
 - Quantify function list 75
- Windows, displaying UNIX coverage data 64
- write-export-file 83, 84
- write-summary-file 83, 84

Y

- yellow memory color 51

