PGI[®] Server 7.1 PGI[®] Workstation 7.1

Installation & Release Notes

The Portland Group™ STMicroelectronics, Inc Two Centerpointe Drive Lake Oswego, OR 97035 www.pgroup.com While every precaution has been taken in the preparation of this document, The Portland GroupTM (PGI®), a wholly-owned subsidiary of STMicroelectronics, Inc., makes no warranty for the use of its products and assumes no responsibility for any errors that may appear, or for damages resulting from the use of the information contained herein. STMicroelectronics, Inc. retains the right to make changes to this information at any time, without notice. The software described in this document is distributed under license from STMicroelectronics, Inc. and may be used or copied only in accordance with the terms of the license agreement. No part of this document may be reproduced or transmitted in any form or by any means, for any purpose other than the purchaser's personal use without the express written permission of STMicroelectronics, Inc.

Many of the designations used by manufacturers and sellers to distinguish their products are claimed as trademarks. Where those designations appear in this manual, STMicroelectronics was aware of a trademark claim. The designations have been printed in caps or initial caps.

PGF95, *PGF90* and *PGC++* are trademarks and *PGI*, *PGHPF*, *PGF77*, *PGCC*, *PGPROF*, and *PGDBG* are registered trademarks of STMicroelectronics, Inc. *Other brands and names are the property of their respective owners.

PGI Server 7.1 / PGI Workstation 7.1 Installation & Release Notes Copyright © 2007

The Portland Group[™] STMicroelectronics, Inc. - All rights reserved. Printed in the United States of America

First Printing:	Release 7.1-1, October, 2007
Second Printing:	Release 7.1-2, November, 2007

Technical support: http://www.pgroup.com/support

Table of Contents

1	PGI R	ELEASE 7.1 INTRODUCTION	1
	1.1 Pr	RODUCT OVERVIEW	1
	1.2 TI	ERMS AND DEFINITIONS	2
2	PGI	RELEASE 7.1 INSTALLATION NOTES	7
	2.1 IN	TRODUCTION	7
	2.2 LI	CENSING	8
	2.2.1 P	GI Workstation Licensing	9
	2.2.2 P	GI Server Licensing	9
	2.2.3 T	rial Licensing Key Constraints	10
	2.2.4 L	icense Keys and System Configurations	10
	2.3 IN	ISTALLING ON LINUX	11
	2.3.1	Preparing to Install on Linux	11
	2.3.2	Installation Steps for Linux	13
	2.3.3	End-user Environment Settings on Linux	20
	2.4 IN	STALLING FLEXLM ON LINUX	22
	2.5 IN	STALLING ON WINDOWS	26
	2.5.1	Preparing to Install on Windows	26
	2.5.2	Installation Steps for Windows	28
	2.5.3	Customizing the Command Window	31
	2.5.4	PGI Workstation Directory Structure	31
	2.5.5	Using LM LICENSE FILE	33
	2.5.6	Common Windows Installation Problems	34
	2.6 IN	STALLING ON APPLE MAC OS X	35
	2.6.1	Preparing to Install on Apple Mac OS X	36
	2.6.2	Installation Steps for Mac OS	38
	2.6.3	End-user Environment Settings on Mac OS	41

3	PGI RELEASE 7.1 RELEASE NOTES	43
	3.1 PGI RELEASE 7.1 CONTENTS	44
	3.2 SUPPORTED SYSTEMS	44
	3.2.1 Supported Processors	44
	3.2.2 Supported Operating Systems	46
	3.3 NEW OR MODIFIED COMPILER FEATURES	48
	3.4 COMPILER OPTIONS	50
	3.4.1 Getting Started	50
	3.4.1.1 Using – fast, – fastsse, and Other Performance-Enhancing	
	Options	50
	3.4.2 New or Modified Compiler Options	51
	3.5 PGDBG NEW AND MODIFIED FEATURES	55
	3.6 PGPROF NEW AND MODIFIED FEATURES	55
	3.7 RUNNING AN MPICH PROGRAM ON LINUX	56
	3.8 USING THE PGI WINDOWS CDK WITH	
	MICROSOFT COMPUTE CLUSTER SERVER	56
	3.8.1 Build MPI Applications with MSMPI	57
	3.8.2 Debug Cluster Applications that Generate MPI Profile Data	57
	3.9 PGI WORKSTATION 7.1 FOR WINDOWS	59
	3.9.1 The Windows Command Environment	59
	3.9.2 MKS Toolkit Compatibility	60
	3.9.3 Using Shared object files in SFU and SUA	60
	3.10 PGI WORKSTATION 7.1 FOR MAC OS X	62
	3.10.1 Mac OS X Debugging Requirements	62
	3.11 GENERATING PGI UNIFIED BINARIES	63
	3.11.1 Unified Binary Command-line Switches	63
	3.11.2 Unified Binary Directives and Pragmas	64
	3.12 STATIC AND DYNAMIC LINKING ON WINDOWS	64
	3.12.1 –Bdynamic	65
	3.12.2 –Bstatic	65
	3.13 USING ENVIRONMENT MODULES	65
	3.14 THE REDIST DIRECTORIES	66
	3.14.1 PGI Redistributables	67
	3.14.2 Microsoft Redistributables	67
	3.15 CUSTOMIZING WITH SITERC AND USER RC FILES	67
	3.16 KNOWN LIMITATIONS	69
	3.17 CORRECTIONS	72
	<i>3.17.1 Corrections in 7.1-2</i>	73
	<i>3.17.2 Corrections in 7.1-1</i>	73
4	CONTACT INFORMATION AND DOCUMENTATION	79

1

PGI Release 7.1 Introduction

Welcome to Release 7.1 of *PGI Workstation* and *PGI Server*, a set of Fortran, C, and C++ compilers and development tools for 32-bit and 64-bit x86-compatible processor-based workstations and servers running versions of the Linux, Windows, and Mac OS operating systems.

All workstation-class compilers and tools products from The Portland Group (*PGHPF Workstation*, for example) are subsets of the *PGI Workstation Complete* product. These workstation-class products provide a node-locked single-user license, meaning one user at a time can compile on the one system on which the *PGI Workstation* compilers and tools are installed.

PGI Server products are offered in configurations identical to the workstationclass products, but provide network-floating multi-user licenses. This means that two or more users can use the *PGI* compilers and tools concurrently on any compatible system networked to the system on which the *PGI Server* compilers are installed.

These release notes apply to all workstation-class and server-class compiler products from The Portland Group.

1.1 Product Overview

Release 7.1 of PGI Workstation and PGI Server includes the following components:

- PGF95 OpenMP* and auto-parallelizing Fortran 90/95 compiler.
- PGF77 OpenMP and auto-parallelizing FORTRAN 77 compiler.
- *PGHPF* data parallel High Performance Fortran compiler. NOTE: *PGHPF is not supported on Windows platforms*.

- *PGCC* OpenMP and auto-parallelizing ANSI C99 and K&R *C* compiler.
- *PGC*++ OpenMP and auto-parallelizing ANSI *C*++ compiler.
- PGPROF graphical MPI/OpenMP/multi-thread performance profiler.
- *PGDBG* graphical MPI/OpenMP/multi-thread symbolic debugger
- *MPICH MPI libraries, version 1.2.7,* for both 32-bit and 64-bit development environments (Linux only)
- Online documentation in PDF, HTML and man page formats.
- A UNIX*-like shell environment for *Win32* and *Win64* platforms.

Depending on the product configuration you purchased, you may not have licensed all of the above components.

The MPI profiler and debugger included with *PGI Workstation* are limited to processes on a single node. *PGI Workstation* can be installed on a single computer, and that computer can be used to develop, debug, and profile MPI applications. The *PGI CDK Cluster Development Kit* supports general development on clusters.

1.2 Terms and Definitions

Following are definitions of terms used in the context of these release notes.

AMD64 – a 64-bit processor from AMD designed to be binary compatible with 32-bit *x86* processors, and incorporating new features such as additional registers and 64-bit addressing support for improved performance and greatly increased memory range. This includes the AMD* Athlon64*, AMD Opteron* and AMD Turion* processors.

Barcelona – In this document the Quad-Core AMD Opteron(TM) Processor (i.e. Opteron Rev x10) is referred to as Barcelona.

DLL – a dynamic linked library on *Win32* or *Win64* platforms of the form *xxx.dll* containing objects that are dynamically linked into a program at the time of execution.

driver – the compiler *driver* controls the compiler, linker, and assembler, and adds objects and libraries to create an executable. The *-dryrun* option illustrates operation of the driver. pgf77, pgf95, pghpf, pgcc, pgCC (Linux), and pgcpp are drivers for the PGI compilers. A pgf90 driver is retained for compatibility with existing makefiles, even though pgf90 and pgf95 drivers are identical.

Dual-, Quad-, or Multi-core – some x64 CPUs incorporate two or four complete processor cores (functional units, registers, level 1 cache, level 2 cache, etc) on a single silicon die. These are referred to as Dual-core or Quad-core (in general, Multi-core) processors. For purposes of OpenMP or auto-parallel threads, or MPI process parallelism, these cores function as distinct processors. However, the processing cores are on a single chip occupying a single socket on the system motherboard. In PGI 7.1, there are no longer software licensing limits on OpenMP threads for Multi-core.

EM64T – a 64-bit IA32 processor with *Extended Memory 64-bit Technology* extensions designed to be binary compatible with AMD64 processors. This includes Intel Pentium 4, Intel Xeon, and Intel Core 2 processors.

FLEXIm – The flexible license management software from Macrovision.

Hyperthreading (HT) – some IA32 CPUs incorporate extra registers that allow 2 threads to run on a single CPU with improved performance for some tasks. This is called *hyperthreading* and abbreviated *HT*. Some *linux86* and *linux86*-64 environments treat IA32 CPUs with HT as though there were a 2^{nd} pseudo CPU, even though there is only one physical CPU. Unless the Linux kernel is *hyperthread-aware*, the second thread of an *OpenMP* program will be assigned to the *pseudo* CPU, rather than a real second physical processor (if one exists in the system). *OpenMP* Programs can run very slowly if the second thread is not properly assigned.

IA32 – an Intel Architecture 32-bit processor designed to be binary compatible with x86 processors, but incorporating new features such as streaming SIMD extensions (SSE) for improved performance. This includes the Intel Pentium* 4 and Intel Xeon* processors. For simplicity, these release notes refer to x86 and IA32 processors collectively as 32-bit x86 processors.

Large Arrays – arrays with aggregate size larger than 2GB, which require 64bit index arithmetic for accesses to elements of arrays. Program units that use *Large Arrays* must be compiled using –*mcmodel=medium*. If –*Mlarge_arrays* is specified and –*mcmodel=medium* is not specified, the default small memory model is used, and all index arithmetic is performed in 64-bits. This can be a useful mode of execution for certain existing 64-bit applications that use the small memory model but allocate and manage a single contiguous data space larger than 2GB.

linux86 – 32-bit Linux operating system running on an *x86*, *AMD64* or *EM64T* processor-based system, with 32-bit GNU tools, utilities and libraries used by the PGI compilers to assemble and link for 32-bit execution.

linux86-64 – 64-bit Linux operating system running on an *AMD64* or *EM64T* processor-based system, with 64-bit and 32-bit GNU tools, utilities and libraries used by the PGI compilers to assemble and link for execution in either *linux86* or *linux86-64* environments. The 32-bit development tools and execution environment under *linux86-64* are considered a cross-development environment for *x86* processor-based applications.

Mac OS – collectively, all *osx86* and *osx86-64* platforms supported by the PGI compilers.

-mcmodel=small – compiler/linker switch to produce *small memory model* format objects/executables in which both code (*.text*) and data (*.bss*) sections are limited to less than 2GB. This switch is the default and only possible format for *linux86* 32-bit executables. This switch is the default format for *linux86-64* executables. Maximum address offset range is 32-bits, and total memory used for OS+Code+Data must be less than 2GB.

-mcmodel=medium – compiler/linker switch to produce medium memory model format objects/executables in which code sections are limited to less than 2GB, but data sections can be greater than 2GB. This option is supported only in *linux86-64* environments. It must be used to compile any program unit that will be linked in to a 64-bit executable that will use aggregate data sets larger than 2GB and will access data requiring address offsets greater than 2GB. This option must be used to *link* any 64-bit executable that will use aggregate data sets greater than 2GB in size. Executables linked using -mcmodel=medium can incorporate objects compiled using -mcmodel=small as long as the small objects are from a shared library.

MPI –MPI is a language-independent communications protocol used to program parallel computers.

MPICH – MPICH is a freely available, portable implementation of MPI, a standard for message-passing for distributed-memory applications used in parallel computing.

Network Installation – A term that applies to installing software on a shared file system available to many machines. Typically this type of installation allows multiple users to access and run the software – up to the number of licenses associated with the software.

NUMA – *Non-Uniform Memory Access.* A type of multi-processor system architecture in which the memory latency from a given processor to a given portion of memory can vary, resulting in the possibility for compiler or programming optimizations to ensure frequently accessed data is "close" to a given processor as determined by memory latency.

osx86 – 32-bit Apple Mac OS Operating Systems running on an x86 Core 2 or Core 2 Duo processor-based system with the 32-bit Apple and GNU tools, utilities, and libraries used by the PGI compilers to assemble and link for 32-bit execution.

osx86-64 - 64-bit Apple Mac OS Operating Systems running on an x64 Core 2 Duo processor-based system with the 64-bit and 32-bit Apple and GNU tools, utilities, and libraries used by the PGI compilers to assemble and link for either 64- or 32-bit execution.

SFU – *Windows Services for Unix* is the precursor to *SUA*. *SFU* supports 32-bit applications on *Windows 2000*, *Windows Server 2003*, and *XP*.

Shared library – a Linux library of the form *libxxx.so* containing objects that are dynamically linked into a program at the time of execution.

SSE - collectively, all SSE extensions supported by the PGI compilers.

SSE1 – 32-bit IEEE 754 FPU and associated *streaming SIMD extensions* (SSE) instructions on Pentium III, AthlonXP* and later 32-bit *x86*, *AMD64* and *EM64T* compatible CPUs, enabling scalar and packed vector arithmetic on single-precision floating-point data.

SSE2 - 64-bit IEEE 754 FPU and associated SSE instructions on P4/Xeon and later 32-bit *x86*, *AMD64* and *EM64T* compatible CPUs. SSE2 enables scalar and packed vector arithmetic on double-precision floating-point data.

SSE3 – additional 32-bit and 64-bit SSE instructions to enable more efficient support of arithmetic on complex floating-point data on 32-bit *x86*, *AMD64* and *EM64T* compatible CPUs with so-called *Prescott New Instructions* (PNI), such as Intel IA32 processors with EM64T extensions and newer generation (Revision E and beyond) AMD64 processors.

SSE4A and ABM – AMD Instruction Set enhancements for the Quad-Core AMD Opteron Processor. Support for these instructions is enabled by the -tp barcelona or -tp barcelona-64 switch.

SSSE3 – an extension of the SSE3 instruction set found on the Intel Core 2.

Static linking – a method of linking:

- On Linux, use *-Bstatic* to ensure all objects are included in a generated executable at link time. Static linking causes objects from static library archives of the form libxxx.a to be linked in to your executable, rather than dynamically linking the corresponding libxxx.so shared library.
- On Windows, the Windows linker links statically or dynamically depending on whether the libraries on the link-line are DLL import libraries or static libraries. By default, the static PGI libraries are included on the link line. To link with DLL versions of the PGI libraries instead of static libraries, compile and link with the *–Bdynamic option*.

SUA – the Subsystem for Unix-based Applications is a source-compatible subsystem for compiling and running 32- and 64-bit UNIX-based applications on a computer running Windows. *SUA* is supported on *Windows 2003 R2* and *Vista*. *SUA* is bundled with Windows; however, the full *SUA Utilities SDK* must be installed to link programs.

Win32 – any of the 32-bit Microsoft* Windows* Operating Systems (XP/2000/Server 2003) running on an *x86*, *AMD64* or *EM64T* processor-based system. On these targets, the PGI compiler products include additional Microsoft tools and libraries needed to build executables for 32-bit Windows systems.

Win64 – any of the 64-bit Microsoft Windows Operating Systems (XP Professional / Windows Server 2003 x64 Editions) running on an x64 processor-based system. On these targets, the PGI compiler products include additional Microsoft tools and libraries needed to build executables for either Win32 or Win64 environments.

Windows – collectively, all *Win32* and *Win64* platforms supported by the PGI compilers.

x64 – collectively, all AMD64 and EM64T processors supported by the PGI compilers.

x86 – a processor designed to be binary compatible with i386/i486 and previous generation processors from Intel* Corporation. Used to refer collectively to such processors up to and including 32-bit variants.

*x*87 – 80-bit IEEE stack-based floating-point unit (FPU) and associated instructions on *x*86–compatible CPUs.

PGI Release 7.1 Installation Notes

2.1 Introduction

2

This section provides an overview of the steps required to successfully install *PGI Workstation* or *PGI Server*. The remainder of this chapter provides the details of each of the steps. Specifically, section 2.2 describes licensing and how to successfully generate either trial or permanent license keys. Section 2.3 describes how to install *PGI Workstation* or *PGI Server* in a generic manner on Linux. Section 2.4 describes how to install and run a FLEXIm license daemon on Linux. Section 2.5 describes how to install *PGI Workstation* on Windows systems and how to install and run a FLEXIm license daemon on Windows. Section 2.6 describes how to install *PGI Workstation* on an Apple system.

Before you begin the installation, it is essential that you understand the flow of the installation process. There are basically three stages of the process:

- **Prepare to install** verifying that you have all the required information and software.
- Install the software installing the software appropriate for your operating system.
- License the software installing the licenses and starting the license server.

The following illustration provides a high-level overview of the installation process.



For more complete information on these steps and the specific actions to take for your operating system, refer to the following sections.

2.2 Licensing

The PGI compilers and tools are license-managed. There are two types of license keys: trial and permanent. This section describes the generation of these license keys.

Note. You must install the software before you obtain your license because the licensing request prompts you for information that is generated during the software installation.

To generate these keys, you use your personalized account on the *http://www.pgroup.com* web page.

- **Permanent keys:** When you purchase a permanent license, the e-mail order acknowledgement you receive includes complete instructions for logging on to the pgroup.com web page and generating permanent license keys.
- **Trial keys:** When you register for a trial, you should generate trial keys using the web page: http://www.pgroup.com/license/trial.

Note. The preview release of *PGI Workstation for Mac OS* has a time-limited test license. No license keys are required for Mac OS during the preview period.

Now let's look at specific licensing issues.

2.2.1 PGI Workstation Licensing

PGI Workstation is licensed to a single system.

On Linux, there are two licensing options for PGI Workstation: PGI-style licensing and FLEXIm-style licensing.

- **PGI-style licensing** for PGI Workstation allows a named user to run as many simultaneous copies of the compiler and tools as desired, but usage is restricted to a pre-specified username.
- **FLEXIm-style licensing** allows any user of the system to use PGI Workstation; however, only a single copy of a compiler or tool is allowed to run at any given time.

On Mac OS and Windows, including SUA and SFU, only FLEXIm-style licensing is supported.

2.2.2 PGI Server Licensing

PGI Server supports multi-user, network-floating licenses. FLEXIm-style licensing is required for *PGI Server*. Multiple users can use the PGI compilers simultaneously from multiple systems on a network when those systems have a properly configured version of PGI Server installed.

On Linux, *PGI Server* may be installed locally on each machine on a network or may be installed once on a shared file system available to each machine. If you select the second method, a *network install*, adding another machine to the group running the compilers is a much simpler process in this release; you

adjust to the unique characteristics of the newly added system with a customization script that must be executed on each machine in the group.

Note: On Windows and Mac OS, *PGI Server* must be installed locally on each machine.

If you require FLEXIm-style licensing, you must request FLEXIm-style license keys when generating permanent keys. First install the PGI compilers and tools according to the instructions in the following sections. You must then install and configure the FLEXIm license management software according to the instructions in sections 2.3 or 2.4. These sections describe how to configure license daemons for Linux and Windows, respectively, including installation of the license daemon and proper initialization of the LM_LICENSE_FILE environment variable.

2.2.3 Trial Licensing Key Constraints

NOTE

At the conclusion of the trial period, the PGI compilers and tools and any executable files generated prior to the installation of permanent license keys will cease to function. Any executables, object files, or libraries created using the PGI compilers with a trial key must be recompiled with permanent license keys in place.

2.2.4 License Keys and System Configurations

Executable files generated with permanent license keys in place are unconstrained, and will run on any compatible system regardless of whether the PGI compilers are installed.

If you change the configuration of your system by adding or removing hardware, your license key may become invalid. Please contact The Portland Group if you expect to reconfigure your system to ensure that you do not temporarily lose the use of the PGI compilers and tools.

For the first 60 days after your purchase, you may send technical questions about these products to the e-mail address *trs@pgroup.com*. If you have purchased a PGI Software Subscription, you have access to e-mail support for an additional 12 months and you are notified by e-mail when maintenance

releases occur and are available for electronic download and installation. Phone support is not currently available. Contact us at *sales@pgroup.com* if you would like information regarding the subscription service for the PGI products you have purchased.

The following sections describe how to install *PGI Workstation* or *PGI Server*.

2.3 Installing on Linux

This section describes how to install *PGI Workstation* or *PGI Server* on a Linux system. It covers local and network installs and is applicable to permanent or trial installations.

- For installations on 32-bit x86 systems, the PGI installation script installs only the linux86 versions of the PGI compilers and tools.
- For installations on 64-bit x64 systems running a linux86-64 execution and development environment, the PGI installation script installs the linux86-64 version of the PGI compilers and tools. Further, if the 32-bit gcc development package is already installed on the system, the 32-bit linux86 tools are also installed on a 64-bit x64 system.

2.3.1 Preparing to Install on Linux

In preparation for installing *PGI Workstation* or *PGI Server* on Linux, first review the following overview of the Linux installation process.



Note. For Linux installations, **each user** must also set their environment variables to properly access the software, as described in 2.3.3 End-user Environment Settings on Linux on page 20.

The 32-bit and 64-bit compilers, tools, and supporting components have the same command names. Further, the environment you target by default (*linux86-64* or *linux86*) depends on the version of the compiler that comes first in your path settings.

In a traditional local installation, the default installation base directory is /opt/pgi. If you choose to perform a network install, you should specify a shared file system for the installation base directory. You must also specify a second directory name that is local to each of the systems where the PGI compilers and tools are used. This local directory will contain the libraries to use when compiling and running on that machine. This approach allows a network installation to support a network of machines that run different versions of Linux.

To prepare for the installation:

- Locate the email that contains your Order Acknowledgement. This email contains your username and password for the web page as well as the PIN number associated with your licensing file.
- Bring up a shell command window on your system. The installation instructions assume you are using csh, sh, ksh, bash, or some compatible shell. If you are using a shell that is not compatible with one of these shells, appropriate modifications are necessary when setting environment variables.
- Verify you have enough free disk space. On the linux86 platform, PGI installation requires 250 MB of free disk space. On the linux86-64 platform, PGI installation requires 1.4 GB of free disk space.

2.3.2 Installation Steps for Linux

Follow these instructions to install the software:

Step 1. If you received this software on a CD-ROM, please skip to step 2. If you downloaded the software from http://www.pgroup.com or another electronic distribution site, then in the instructions that follow, replace <tarfile> with the name of the file that was downloaded.

Note. The PGI products cannot be installed into the same directory where the tar file is unpacked. Use the following command sequence to unpack the tar file in a temporary directory before installation:

```
% mkdir /tmp/pgi
% mv <tarfile>.tar.gz /tmp/pgi
% cd /tmp/pgi
% tar xpfz <tarfile>.tar.gz
```

Step 2. Run Install Script.

The install script *must* be run to properly install the software.

The install script lists the products that are available on the CD-ROM or in the download package. You will be asked which products should be installed, whether to perform a traditional local installation or a network installation, and where to place the installation directory.

After the software is installed, the install script performs some system-specific customization and then initializes the licensing, as described in step 3.

• If you downloaded the software from the Internet, execute the following script in the directory where you unpacked the tar file:

```
% ./install
```

If you are installing from a CD-ROM, issue the following command:
 % /mnt/cdrom/install

Note. For a network installation, you are asked for a common local directory. This local directory will be created once on each system utilizing the network installation; further, it must be created on each system *before* adding that system to the network using the compilers.

Note. If you have difficulty running this script, especially on a Slackware Linux system, check the permissions on /dev/null. Permissions should be set to "crw-rw-rw-". If necessary, reset permissions to this value; to do this, super-user permissions are required.

Note. Some systems use a CD-ROM volume manager that may insert an additional directory in the above pathname, and require the longer pathname. For example, you may need to execute a script like this:

% /cdrom/pgisoft/install

If you are not sure how to access the CD-ROM drive, check with your system administrator.

Step 3. Get and Load License.

All of the PGI compilers and tools are license-managed. *PGI Workstation* products that are PGI-style licensed and limited to a single user have no need to run a FLEXIm license daemon.

To obtain licensing information, you need the following information:

- The username and password required to connect to the website.
 - For permanent license keys, this information is in your order acknowledgement email.
 - For trial license keys, this is the username and password you used to download the installation software from the web site.
- The FLEXIm hostid and hostname of the computer on which the software is installed, provided by the installation script. You can also obtain your FLEXIm hostid by using the following command after you have installed the products and initialized the environment variables:

% lmutil lmhostid

If you want the *PGI Workstation* compilers to be usable by any one user rather than locked to a specific username, or if you are installing a multi-user *PGI Server* product, you must use FLEXIm. To use FLEXIm, you must specifically request FLEXIm-style keys when generating license keys on the PGI web page at *http://www.pgroup.com/support/keylogin.htm*.

If you have purchased the compilers and tools that you are installing, you received an order acknowledgement e-mail with instructions on how to generate your license keys through the *pgroup.com* web page. The installation script asks for your real name, your username, and your email address and prints a message similar to this:

To obtain an evaluation license, go to: https://www.pgroup.com/license/evaluation.php and use your web-user access codes (email address and password) and the information below to generate a trial license.

For a permanent license, please read the order acknowledgement that you received. Connect to

https://www.pgroup.com/support/keylogin.htm with the username and password provided in the order acknowledgement.

```
FLEX1m hostid: <your host id>
Hostname: <your host name>
Installation: /opt/pgi
PGI Release: 7.1-2
```

For retrieval at a later time, the preceding message is also saved to the file /opt/pgi/license.info, where /opt/pgi is the installation directory.

Once you have obtained your trial or permanent license keys using your personalized account on the *pgroup.com* web page, place them in the file /opt/pgi/license.dat, or substitute the appropriate installation directory path if you have not installed in the default /opt/pgi directory. For more information, refer to Step 2 of Installing FLEXIm on Linux on page 22.

Step 4. Review Documentation.

You can view the online HTML and PDF documentation using any web browser by opening the file:

or

/opt/pgi/linux86/7.1/doc/index.htm

/opt/pgi/linux86-64/7.1/doc/index.htm

You may want to bookmark this location for easy future reference to the online manuals.

Step 5. Make Software Accessible

With either the trial or permanent license file in place, execute the following commands to make the products accessible.

For x64 *linux86-64*:

To use the x64 *linux86-64* version of the compilers and tools, execute the following commands.

Note. In these commands, the installation directory is the default: /opt/pgi.

For csh:

```
% set path = (/opt/pgi/linux86-64/7.1/bin $path)
% setenv MANPATH "$MANPATH":/opt/pgi/linux86-64/7.1/man
```

For bash, sh or ksh:

```
$ PATH=/opt/pgi/linux86-64/7.1/bin:$PATH
```

```
$ export PATH
```

```
$ MANPATH=$MANPATH:/opt/pgi/linux86-64/7.1/man
```

```
$ export MANPATH
```

For linux86:

To use only the *linux86* version of the compilers and tools, or to target *linux86* as the default, use a setup similar to the previous one, changing the path settings as illustrated in the following commands.

For csh:

```
% set path = (/opt/pgi/linux86/7.1/bin $path)
% setenv MANPATH "$MANPATH":/opt/pgi/linux86/7.1/man
```

For bash, sh or ksh:

```
$ PATH=/opt/pgi/linux86/7.1/bin:$PATH
$ export PATH
$ MANPATH=$MANPATH:/opt/pgi/linux86/7.1/man
$ export MANPATH
```

To ensure access to the PGI compilers and tools for future logins, we recommend all users of the PGI products add these commands to their startup files.

Step 6. Verify the Release Number.

To verify the release number of the products you have installed, use the -V option on any of the compiler commands, as illustrated in the following examples. If you use -v instead, you also see the sequence of steps the compiler will use to compile and link programs for execution on your system.

•	For Fortran 77, use:	pgf77 -V x.f
•	For Fortran 95, use:	pgf95 -V x.f
•	For HPF, use:	pghpf -V x.f
•	For C++, use:	pgCC -V x.c
•	For ANSI C, use:	pgcc -V x.c

Note. To successfully execute these commands, the files x.f or x.c need not exist.

Step 7. Complete Network Installation Tasks.

Skip this step if you are not using a network installation.

For a network installation, you must run the local installation script on each system on the network where the compilers and tools will be available for use.

If your installation base directory is /opt/pgi and the common local directory is /usr/pgi/shared/7.1, run the following commands on each system on the network.

```
/opt/pgi/linux86/7.1-2/bin/makelocalrc \
    -x /opt/pgi/linux86/7.1-2 \
    -net /usr/pgi/shared/7.1
```

```
/opt/pgi/linux86-64/7.1-2/bin/makelocalrc \
    -x /opt/pgi/linux86-64/7.1-2 \
    -net /usr/pgi/shared/7.1
```

These commands create a system-dependent file *localrc.machinename* in both the /opt/pgi/linux86/7.1-2/bin directory and in /opt/pgi/linux86-64/7.1-2/bin. The commands also create the following three directories containing libraries and shared objects specific to the operating system and system libraries on that machine:

```
/usr/pgi/shared/7.1/lib, /usr/pgi/shared/7.1/liblf, and
/usr/pgi/shared/7.1/lib64.
```

Note. The makelocalrc command does allow the flexibility to have local directories that have different names on different machines. However, using the same directory on different machines allows users to easily move executables that use PGI-supplied shared libraries between systems.

If you specify /opt/pgi as the base directory for installation, the following directory structure is created by the PGI installation script:

Name of directory	Contents
/opt/pgi/linux86/7.1/bin	<i>linux86</i> 32-bit compilers & tools
/opt/pgi/linux86/7.1/lib	<i>linux86</i> 32-bit runtime libraries
/opt/pgi/linux86/7.1/liblf	<i>linux86</i> 32-bit large-file support libs (used by – <i>Mlfs</i>)
/opt/pgi/linux86/7.1/include	<i>linux86</i> 32-bit header files
/opt/pgi/linux86-64/7.1/bin	<i>linux86-64</i> compilers & tools
/opt/pgi/linux86-64/7.1/lib	<i>linux86-64 –mcmodel=small</i> libs

Name of directory	Contents	
/opt/pgi/linux86-64/7.1/libso	<i>linux86-64 –fpic</i> shared libraries for <i>–mcmodel=medium</i> development	
/opt/pgi/linux86-64/7.1/include	<i>linux86-64</i> header files	
/opt/pgi/linux86/7.1/REDIST /opt/pgi/linux86-64/7.1/REDIST	Re-distributable runtime libraries	
/opt/pgi/linux86/7.1/EXAMPLES /opt/pgi/linux86-64/7.1/EXAMPLES	Compiler examples	
/opt/pgi/linux86/7.1/doc /opt/pgi/linux86-64/7.1/doc	Documentation	
/opt/pgi/linux86/7.1/man /opt/pgi/linux86-64/7.1/man	UNIX-style man pages	
/opt/pgi/linux86/7.1/jre /opt/pgi/linux86-64/7.1/jre	JAVA environment for <i>PGDBG</i> and <i>PGPROF</i> graphical user interfaces	
/opt/pgi/linux86/7.1/src /opt/pgi/linux86-64/7.1/src	Fortran 90 source files for included modules.	
/opt/pgi/linux86/7.1/mpi/mpich /opt/pgi/linux86-64/7.1/mpi/mpich	MPICH tools and libraries.	

Additionally, a network install creates the following directories:

Name of directory	Contents	
/opt/pgi/linux86/7.1/lib-linux86-g	<i>linux86</i> 32-bit <i>libpgc</i> library dependent on the version of <i>glibc</i> installed on each machine	
/opt/pgi/linux86/7.1/include-g	<i>linux86</i> 32-bit header files dependent on the version of glibc or gcc installed on each machine	
/opt/pgi/linux86-64/7.1/include-g	<i>linux86-64</i> 64-bit header files dependent on the version of <i>glibc</i> or <i>gcc</i> installed on each machine	

2.3.3 End-user Environment Settings on Linux

Now that you have installed the compilers in, for example, /opt/pgi, you must initialize your environment to use the compilers successfully. Assume the license file is in /opt/pgi/license.dat, and the *lmgrd* license manager is running.

Note. Each user must issue the following sequence of commands to initialize the shell environment before using the PGI compilers and tools.

32-bit Compiler Commands

The following commands make the 32-bit compilers the default.

• In csh, use these commands:

```
% setenv PGI /opt/pgi
% setenv MANPATH "$MANPATH":$PGI/linux86/7.1/man
% setenv LM_LICENSE_FILE $PGI/license.dat
% set path = ($PGI/linux86/7.1/bin $path)
```

To access the MPICH commands, execute this command line:

% set path = (\$PGI/linux86/7.1/mpi/mpich/bin \$path)

To access the MPICH man pages, execute this command line:

- % setenv MANPATH \
 "\$MANPATH":\$PGI/linux86/7.1/mpi/mpich/man
- In bash, sh or ksh, use these commands:
 - \$ PGI=/opt/pgi; export PGI
 - \$ MANPATH=\$MANPATH:\$PGI/linux86/7.1/man
 - \$ export MANPATH
 - \$ LM_LICENSE_FILE=\$PGI/license.dat
 - \$ export LM_LICENSE_FILE
 - \$ PATH=\$PGI/linux86/7.1/bin:\$PATH
 - \$ export PATH

To access the MPICH commands, execute this command line:

\$ PATH=\$PGI/linux86/7.1/mpi/mpich/bin:\$PATH

To access the MPICH man pages, execute this command line:

```
$ MANPATH=\
$MANPATH:$PGI/linux86/7.1/mpi/mpich/man
```

64-bit Compiler Commands

The following commands make the 64-bit compilers the default.

```
• In csh, use these commands:

% setenv PGI /opt/pgi

% setenv MANPATH "$MANPATH":$PGI/linux86-64/7.1/man

% setenv LM_LICENSE_FILE $PGI/license.dat

% set path = ($PGI/linux86-64/7.1/bin $path)
```

To access the MPICH commands, execute this command line:

```
% set path = ($PGI/linux86-64/7.1/mpi/mpich/bin $path)
```

To access the MPICH man pages, execute this command line:

```
% setenv MANPATH \
    "$MANPATH":$PGI/linux86-64/7.1/mpi/mpich/man

• In bash, sh or ksh, use these commands:
    $ PGI=/opt/pgi; export PGI
    $ MANPATH=$MANPATH:$PGI/linux86-64/7.1/man
    $ export MANPATH
    $ LM_LICENSE_FILE=$PGI/license.dat
    $ export LM_LICENSE_FILE
    $ PATH=$PGI/linux86-64/7.1/bin:$PATH
    $ export PATH
```

To access the MPICH commands from bash, sh, or ksh:

```
$ PATH=$PGI/linux86-64/7.1/mpi/mpich/bin:$PATH
```

To access the MPICH man pages from bash, sh, or ksh:

```
$ MANPATH=\
$MANPATH:$PGI/linux86-64/7.1/mpi/mpich/man
```

2.4 Installing FLEX1m on Linux

If you want the *PGI Workstation* to be usable by any one user, rather than locked to a specific *username*, or if you are installing a multi-user *PGI Server* product, you must use the FLEXIm software license management system from Macrovision* Software, as outlined in this section.

PGI Workstation trial keys use FLEXIm software license management. Please follow these directions to enable the trial.

IMPORTANT NOTE: Release 7.1 includes a newer version of the Macrovision FLEXIm software. The updated *lmgrd* and *pgroupd* daemons must be used in preference to versions shipped with previous releases of the PGI products. You can co-install Release 7.1 with Release 7.0, 6.X and/or 5.2; and you can use any of these versions of the compilers and tools with the new versions of *lmgrd* and *pgroupd* and a single Release 7.1 license file. Further, if you use this file to start *lmgrd* automatically after a reboot of your system, you must modify your lmgrd.rc file in the /etc/rc.d or /etc/init.d directory to use the new *lmgrd*.

For example, your lmgrd.rc file may look like this one, where <target> is replaced appropriately with linux86 or linux86-64.

```
## Path to master daemon lmgrd
# Commented out previous path to 5.2:
#LMGRD=$PGI/<target>/5.2/bin/lmgrd
LMGRD=$PGI/<target>/7.1/bin/lmgrd
## Command to stop lmgrd
#Commented out previous path to 5.2:
#LMUTIL=$PGI/<target>/5.2/bin/lmutil
LMUTIL=$PGI/<target>/7.1/bin/lmutil
```

For complete details on setup and usage of these files, see Step 4 on page 24.

Step 1. Install the PGI software as described in section 2.2 on page 11.

Step 2. Obtain and Install License Keys.

Once you have obtained trial or permanent FLEXIm-style license keys, as described in Step 3 of section 2.2, place them in a file named license.dat in the installation directory, typically, the /opt/pgi directory. For example, if you have purchased *PGF77 Workstation* for Linux, the license.dat file looks similar to the following:

```
SERVER <hostname> <hostid> 7496
DAEMON pgroupd pgroupd
FEATURE pgf77-linux86 pgroupd 7.100 31-dec-0 1 \
2B9CF0F163159E4ABE32 VENDOR_STRING=107209:16 \
HOSTID=<hostid> ck=49
FEATURE pgprof pgroupd 7.100 31-dec-0 1 \
6BDCE0B12EC19D0909F0 VENDOR_STRING=107209:16 \
HOSTID=<hostid> ck=60
```

In your license file:

- <hostname> and <hostid> must match those you submitted when obtaining your licenses.
- In the feature line component, VENDOR_STRING=107209:16 includes the Product ID Number (PIN) for this installation, in this example 107209. You have a similar unique PIN number for your installation.

Note. Please include your PIN number when sending mail to us regarding technical support for the products you have purchased. This PIN number is also on your order acknowledgement email.

Step 3. Access products and initialize your environment.

When the license file is in place, execute the following commands to make the products accessible and to initialize your environment for use of FLEXIm. These commands assume that you use the default installation directory: /opt/pgi

```
In csh, use these commands:
% setenv PGI /opt/pgi
% setenv LM_LICENSE_FILE \
    "$LM_LICENSE_FILE":/opt/pgi/license.dat
In bash, sh or ksh, use these commands:
$ PGI=/opt/pgi
$ export PGI
$ LM_LICENSE_FILE= \
    $LM_LICENSE_FILE= / opt/pgi/license.dat
$ export LM LICENSE FILE
```

You should add these commands to your startup files to ensure that you have access to the PGI products upon future logins.

Step 4. Start the license manager daemon.

To complete your installation, you must now start the license manager daemon.

Important: If you are evaluating PGI Workstation with trial keys, you are done. You do not need to perform this step.

Note. The following paragraph refers to the shell script template for *linux86*. If you have installed only *linux86-64*, please substitute *linux86-64* for *linux86*.

Edit the shell script template /opt/pgi/linux86/7.1/bin/lmgrd.rc, where /opt/pgi is the default installation directory. If you installed the compilers in a directory other than /opt/pgi, substitute the correct installation directory for '/opt/pgi' on line 3 of the script. Then save the file and exit the editor.

Issue the following command to start the license server and *pgroupd* license daemon running on your system:

```
% lmgrd.rc start
```

If you wish to stop the license server and license daemon at a later time, you can do so with the command:

```
% lmgrd.rc stop
```

To make sure that the license server and *pgroupd* daemon are started each time your system is booted, log in as root, set the PGI environment variable as described in Step 3 on page 23, and then execute the following two commands:

```
% cp /opt/pgi/linux86/7.1/bin/lmgrd.rc/etc/init.d/lmgrd
% ln -s /etc/init.d/lmgrd/etc/rc.d/rc3.d/S90lmgrd
```

Note. There are two values in this example that may be different on your system:

• Your system's default runlevel may be something other than '3', the level used in this example. You can run /sbin/runlevel to check the system's runlevel. If the runlevel on your systems is different, then you must set the correct subdirectory; use your system's runlevel in place of the "3" in the preceding example.

• Your rc files may be in a directory other than the one in the example: /etc/init.d. If the rc files are in a directory such as /etc/rc.d/init.d, then substitute that location in the example.

Some Linux distributions, such as SuSE and Red Hat, include the *chkconfig(8)* utility that manages the runlevel scripts. If your system has this tool and you wish to use it, then run the following commands:

```
% cp /opt/pgi/linux86/7.1/bin/lmgrd.rc /etc/rc.d/init.d/
% chkconfig --add lmgrd.rc
```

The appropriate links will be created in the /etc/rc.d directory hierarchy. For more information on *chkconfig*, please see the manual page.

Installation of your FLEXIm-style licensing of our products for Linux is now complete. For assistance with difficulties related to the installation, send e-mail to *trs@pgroup.com*.

2.5 Installing on Windows

This section describes how to install *PGI Workstation* on a Windows system, including Win64, Win32, SFU, and SUA. It is applicable to both permanent and trial installations.

- For installations on 32-bit x86 systems, the PGI installer installs only the 32-bit versions of the PGI compilers and tools.
- For installations on 64-bit x64 systems running a 64-bit operating system, the PGI installer installs the 64-bit and 32-bit versions of the PGI compilers and tools.

2.5.1 Preparing to Install on Windows

PGI Workstation for Windows and SFU/SUA includes the *Microsoft Open Tools*, the essential tools and libraries required to compile, link, and execute programs on Windows. No additional Microsoft tools or libraries are needed. The *Microsoft Open Tools* includes a subset of the full Microsoft Platform SDK. *PGI Workstation 7.1* can also compile and link against the Microsoft Platform SDK. For information about how to download the Platform SDK, visit *http://msdn.microsoft.com/platformsdk*.

For SFU and SUA installations, the GNU SDK and GNU utilities must be installed. For SFU, at installation time you must select both the "Interix SDK" and "Interix GNU SDK" components from the Custom Installation. For SUA, you must download the *Utilities and Software Development Kit (SDK) for UNIX-based Applications* from *http://www.microsoft.com* and install the Base SDK and GNU SDK.

In preparation for installing *PGI Workstation* or *PGI Server* on Windows, first review the following overview of the Windows installation process.

Note. You must install the software prior to getting the licenses.



2.5.2 Installation Steps for Windows

Once you have prepared for the installation, follow these instructions to install the software:

Step 1. Have the software available and log on as Administrator.

Administrator privileges are required to install PGI Workstation.

If you received this software on a CD-ROM, please skip to step 2. If not, download the software from http://www.pgroup.com or another electronic distribution site.

Step 2. Start the installation.

If you are installing from a CD-ROM, insert the CD-ROM into the CD-ROM drive on the system on which the install is to take place. An installation script is automatically invoked and the installation process begins. Choose the software you wish to install, and follow the directions printed to your screen.

If you have configured your system so that CD-ROM autorun is disabled, run the installation executable from the CD. If you obtained *PGI Workstation* from PGI electronically, execute this file on the target machine.

pgiws-712.exe	32-bit Windows	
pgiwsx64-712.exe	64-bit/32-bit Windows	
pgiws-sfu-712.exe	32-bit SFU	
pgiws-sua-712.exe	32-bit SUA	
pgiwsx64-sua-712.exe	64-bit/32-bit SUA	

The installation executables are:

Step 3. Install the license keys.

The *PGI Workstation* compilers and tools on Windows are license-managed and must use the FLEXIm software license management system from Macrovision Software. This system requires that you possess valid license keys for the licensed product. Further, you must specifically request FLEXIm-style keys when generating license keys at http://www.pgroup.com/support/keylogin.htm.

To obtain licensing information, you need the following information:

- The username and password required to connect to the website.
 - For permanent license keys, this information is in your order acknowledgement email.
 - For trial license keys, this is the username (email address) and password you used to download the installation software from the web site.
- The FLEXIm hostid and hostname of the computer on which the software is installed, provided by the installation script. After you have installed the products, you can also obtain your FLEXIm hostid by opening the PGI workstation window and typing these commands:

PGI\$ **cd \$PGI** PGI\$ **cat license.info**

You see information similar to the following:

For a permanent license, please read the order acknowledgment that you received. Connect to https://www.pgroup.com/support/keylogin.php with the username and password in your order acknowledgment.

```
FLEX1m Host ID: 001234A98765
Installation: C:\Program Files\PGI\
PGI Release: 7.1
```

If you want to obtain an evaluation license, go to:

https://www.pgroup.com/license/evaluation.php and use your web-user access codes (email address and password) and the FLEXIm Host ID to generate a trial license.

For retrieval at a later time, the preceding message is also saved to the file license.info in the installation directory. For example, if you use the default installation directory of C:\Program Files\PGI, then the file location is this:

C:\Program Files\PGI\license.info

Once you have obtained your trial or permanent license keys using your personalized account on the *pgroup.com* web page, place them in the license file: license.dat. In a typical configuration, where C: \ is the system drive and you installed the software using the default location, this file would be found in:

C:\Program Files\PGI\license.dat

- If you have never received license keys from PGI before, replace the license.dat file created during installation with the PGI Workstation keys.
- If your license.dat file already contains keys you have received from PGI, append the PGI Workstation keys to the keys already in this file.

If you are evaluating PGI Workstation with trial keys, skip to Step 5. You do not need to start the license server.

Step 4. Start the PGI License Server.

The FLEXIm license system requires that a license server be running. The installation process creates a Windows Service called PGI License Server. As soon as a valid license.dat file is in place, as described in Step 3, this service can be started.

Important: You do not need to start the license server with trial keys.

The PGI License Server is a Windows Service. Therefore, to start it, do this:

- 1.) Open the Services dialog from the Start menu: (Start | Control Panel | Administrative Tools | Services).
- 2.) Select "PGI License Server".
- 3.) Select "Start."

Note. The PGI License Server service starts automatically on system reboot, provided that the license.dat file contains valid keys.

Step 5. Review Documentation.

You can view the online HTML and PDF documentation using any web browser by opening the file:

```
http://www.pgroup.com/resources/docs.htm
```

You may want to bookmark this location for easy future reference to the online manuals.

Step 6. Verify Release Number.

Verify the release number of the products you have installed. To do so, open PGI Workstation from your desktop and read the first line displayed in the file.

Step 7. Customize PGI workstation.

Optionally, customize the setup, as described in 2.5.3 Customizing the Command Window and in 2.5.5 Using LM_LICENSE_FILE.

2.5.3 Customizing the Command Window

By default, when you double-left-click on the *PGI Workstation* desktop icon, a standard black-background command window appears on your screen. This window is pre-initialized with environment and path settings for use of the *PGI Workstation* compilers and tools. If you prefer different background or text colors, font style, window size, or scrolling capability, you can customize the "shortcut" that creates the *PGI Workstation* command window.

To customize your window, right-click the *PGI Workstation* desktop icon, and select "Properties" from the pop-up menu. In the PGI Workstation Properties dialog box, select the tabs for the features you want to customize, and make the desired modifications.

2.5.4 PGI Workstation Directory Structure

The PGI Workstation default installation directory depends on your platform.

On Win32, the default installation directory is

%SYSTEMDRIVE%\Program Files\PGI\win32\7.1-2\

On Win64 platforms, the default installation directories are

```
%SYSTEMDRIVE%\Program Files\PGI\win64\7.1-2\
%SYSTEMDRIVE%\Program Files (x86)\PGI\win32\7.1-2\
```

In addition to these two product directories, the Microsoft Open Tools and, optionally, Cygwin, are installed here:

```
%SYSTEMDRIVE%\Program Files\PGI\Microsoft Open Tools 8
%SYSTEMDRIVE%\cygwin
```

The *Cygwin* directory is not installed with *PGI Workstation* for *SFU* and *SUA*. Instead, any *SUA* or *SFU* shell can be used. The *PGI Workstation* installation directory structure for SFU and SUA is similar to the Linux directory layout described in Chapter 2. You should also follow instructions in section 2.3.3 End-user Environment Settings on Linux.

The following directory structure will be created during the installation on a Win64 system:

Name of directory	Contents
C:\Program Files\PGI\win64\7.1\bin C:\Program Files (x86)\PGI\win32\7.1\bin	<i>PGI Workstation 7.1</i> compilers and tools binaries
C:\Program Files\PGI\win64\7.1\lib C:\Program Files (x86)\PGI\win32\7.1\lib	<i>PGI Workstation 7.1</i> runtime and support libraries
C:\Program Files\PGI\win64\7.1\include C:\Program Files (x86)\PGI\win32\7.1\include	<i>PGI Workstation 7.1</i> header files
C:\Program Files\PGI\win64\7.1\REDIST C:\Program Files (x86)\PGI\win32\7.1\REDIST	Re-distributable runtime libraries
C:\Program Files\PGI\win64\7.1\doc C:\Program Files (x86)\PGI\win32\7.1\doc	Documentation
C:\Program Files\PGI\win64\7.1\man C:\Program Files (x86)\PGI\win32\7.1\man	Man pages for commands
C:\Program Files\PGI\Microsoft Open Tools 8	Microsoft tools
C:\cygwin	Cygwin tools
The following directory structure will be created during the installation on a Win32 system:

Name of directory	Contents
C:\Program Files\PGI\win32\7.1\bin	<i>PGI Workstation 7.1</i> compilers and tools binaries
C:\Program Files\PGI\win32\7.1\lib	<i>PGI Workstation 7.1</i> runtime and support libraries
C:\Program Files\PGI\win32\7.1\include	PGI Workstation 7.1 header files
C:\Program Files\PGI\win32\7.1\REDIST	Re-distributable runtime libraries
C:\Program Files\PGI\win32\7.1\doc	Documentation
C:\Program Files\PGI\win32\7.1\man	Command Man pages
C:\Program Files\PGI\Microsoft Open Tools 8	Microsoft tools
C:\cygwin	Cygwin tools

2.5.5 Using LM_LICENSE_FILE

The system environment variable LM_LICENSE_FILE is not required by *PGI Workstation* on Windows but you can use it to override the default location that is searched for the license.dat file.

To use the system environment variable LM_LICENSE_FILE, set it to the full path of the license key file. To do this, follow these steps:

- 1. Open the System Properties dialog (Start | Control Panel | System).
- 2. Select the 'Advanced' tab.

- 3. Click the 'Environment Variables' button.
 - If LM_LICENSE_FILE is not already an environment variable, create a new system variable for it. Set its value to the full path, including the name of the file, for the license key file.
 - If LM_LICENSE_FILE already exists as an environment variable, append the path to the license file to the variable's current value using a semi-colon to separate entries.

2.5.6 Common Windows Installation Problems

The most common installation problems on Windows are related to licensing.

To troubleshoot your installation, first check that the license.dat file you are using contains valid license keys. Second, check that the PGI License Server, a Windows Service, has been started.

Typical FLEXIm errors encountered may include the following:

• When using a PGI compiler or tool, a Flexible License Manager dialog appears that states 'LICENSE MANAGER PROBLEM: No such feature exists.'

This message may appear because the license.dat file accessed by the FLEXIm License Manager does not contain valid license keys.

• When using a PGI compiler or tool, a Flexible License Manager dialog appears that states 'LICENSE MANAGER PROBLEM: Cannot connect to license server system.'

This message may appear because the PGI License Server has not been started.

- When starting the PGI License Server, a system message appears that states 'The PGI License Server service on Local Computer started and then stopped. Some services stop automatically if they have no work to do, for example, the Performance Logs and Alerts service.' This message may appear because the license.dat file accessed by the FLEXIm License Manager does not contain valid license keys.
- A message stating 'LICENSE MANAGER PROBLEM: Failed to checkout license' appears. This message may appear because the PGI License Server has not been started.

- By default, on Windows, the license server creates interactive pop-up messages to issue warning and errors. You can use the environment variable FLEXLM_BATCH prevent interactive pop-up windows. To do this, set the environment variable FLEXLM_BATCH to 1.
- On SFU, if the compilers get segmentation faults or produce core dumps, the problem might be related to Windows Data Execution Prevention. The solution to these errors is to modify the system boot.ini file to set /noexecute=AlwaysOff; and then reboot. For more information, refer to this good online resource for SFU: http://www.interopsystems.com.

2.6 Installing on Apple Mac OS X

This section describes how to install *PGI Workstation* on an Apple system. It covers local installs, and is applicable to permanent or trial installations.

Note: PGI Workstation for Mac OS is only supported on Intel Core and Core 2 Duo processors running Mac OS X version 10.4.9. Previous versions of Mac OS may be unstable for 64-bit programs. Using this release requires that Apple Xcode 2.4.1 be installed. Xcode is available from *http://developer.apple.com*.

For installations on 32-bit *x86* systems, the PGI installation process installs only the *osx86* versions of the PGI compilers and tools. For installations on 64-bit *x64* systems running an *osx86-64* execution and development environment, the PGI installation process installs the *osx86-64* version of the PGI compilers and tools. Additionally, if the 32-bit *gcc* development package is already installed on the system, the 32-bit *osx86* tools will be installed on a 64-bit *x64* system.

The 32-bit and 64-bit compilers, tools, and supporting components have the same command names, and the environment you target by default, either *osx86-64* or *osx86*, depends on the version of the compiler that comes first in your path settings.

The default installation base directory is /opt/pgi.

2.6.1 Preparing to Install on Apple Mac OS X

To prepare for the installation:

- Verify you have enough free disk space.
 - On the osx86 platform, PGI installation requires 250 MB of free disk space.
 - On the osx86-64 platform, PGI installation requires 500 MB of free disk space.
- Verify that Xcode 2.4.1 is installed.
 - If you know how to run Xcode, start Xcode and click About Xcode to verify the version is 2.4.1.
 - If you do not know how to run Xcode or are uncertain if it is installed on your system, do the following:
 - 1.) From the Apple Menu, select About This Mac.
 - 2.) Click More Info.
 - 3.) Select System profiler | Software | Applications.
 - 4.) Scroll through the alphabetical list and verify Xcode is in it.
 - 5.) Verify the version is 2.4.1.

Note. PGI Workstation for Mac OS requires the Xcode application, which provides several components of the tool chain, including the system assembler, linker, and run-time libraries. However, PGI Workstation runs in a Terminal, not in the Xcode IDE, and the PGDBG debugger and PGPROF profiler open Java windows.

In preparation for installing *PGI Workstation on Mac OS*, first review the following overview of the Mac OS installation process.



2.6.2 Installation Steps for Mac OS

Once you have prepared for the installation, follow these instructions to install the software:

Step 1. Access the installation package.

- If you received this software on a CD-ROM, place the CD-ROM in the CD drive and wait for the disk to mount of your system.
- If you do not have a CD-ROM, download the software from http://www.pgroup.com or another electronic distribution site. The file you download appears on your system as pgiosx-712.dmg. Open this file to mount it.

Step 2. Install the software.

Click on PGI Workstation.pkg, which is part of the mounted disk. Follow the installation instructions.

- When you see the initial system check dialog, click continue to allow the install script to check that your system has the required components for installing the software, such as Xcode 2.4.1 and gcc.
- While you must select the hard drive on your system for the installation, the install program does not allow you to select an installation directory other than the default one: /opt/pgi.
- After the software is installed, the install script performs some systemspecific customization and then initializes the licensing, as described in step 3.

Step 3. All of the PGI compilers and tools are license-managed.

Step 4. You can view the online HTML and PDF documentation using any web browser by opening the file:

file:/opt/pgi/osx86/7.1/doc/index.htm

or

file:/opt/pgi/osx86-64/7.1/doc/index.htm

You may want to bookmark this location for easy future reference to the online manuals.

Step 5. Make Tools Accessible

To use the compilers and tools, execute the following commands.

For x64 osx86-64:

To use the x64 *osx86-64* version of the compilers and tools, execute the following commands.

Note. In these commands, the installation directory is the default: /opt/pgi. You currently cannot change this directory.

- In bash, zsh, or ksh, use these commands:
 - \$ PATH=/opt/pgi/osx86-64/7.1/bin:\$PATH
 - \$ export PATH
 - \$ MANPATH=\$MANPATH:/opt/pgi/osx86-64/7.1/man
 - \$ export MANPATH
- In csh, use these commands:

```
% set path = (/opt/pgi/osx86-64/7.1/bin $path)
```

```
% setenv MANPATH "$MANPATH":/opt/pgi/osx86-64/7.1/man
```

For osx86:

To install only the *osx86* version of the compilers and tools, or to target *osx86* as the default, use the setup similar to the previous one, changing the path settings:

• In bash, zsh, or ksh, use these commands:

```
$ PATH=/opt/pgi/osx86/7.1/bin:$PATH
$ export PATH
$ MANPATH=$MANPATH:/opt/pgi/osx86/7.1/man
$ export MANPATH
```

• In csh, use these commands:

```
% set path = (/opt/pgi/osx86/7.1/bin $path)
% setenv MANPATH "$MANPATH":/opt/pgi/osx86/7.1/man
```

To ensure they have access to the PGI compilers and tools upon future logins, we recommend that all users of the PGI products add these commands to their startup files.

Step 6 – Verify Release Number.

To verify the release number of the products you have installed, use the -V option on any of the compiler commands, as illustrated in the following examples. If you use -v instead, you also see the sequence of steps the compiler will use to compile and link programs for execution on your system.

- For Fortran 77, use: pgf77 -V x.f
- For Fortran 95, use: pgf95 -V x.f
- For ANSI C, use: pgcc -V x.c

Note. To successfully execute these commands, the files $x \cdot f$ or $x \cdot c$ need not exist.

If you specify /opt/pgi as the base directory for installation, the following directory structure is created during the PGI installation process:

Name of directory	Contents
/opt/pgi/osx86/7.1/bin	osx86 32-bit compilers & tools
/opt/pgi/osx86/7.1/lib	osx86 32-bit runtime libraries
/opt/pgi/osx86/7.1/include	osx86 32-bit header files
/opt/pgi/osx86/7.1/doc	Documentation
/opt/pgi/osx86/7.1/man	UNIX-style man pages
/opt/pgi/osx86/7.1/src	Fortran 90 source files for included modules.
/opt/pgi/osx86-64/7.1/bin	osx86-64 compilers & tools
/opt/pgi/osx86-64/7.1/lib	osx86-64 –mcmodel=small libs
/opt/pgi/osx86-64/7.1/include	osx86-64 header files
/opt/pgi/osx86-64/7.1/doc	Documentation
/opt/pgi/osx86-64/7.1/bin	osx86-64 compilers & tools
/opt/pgi/osx86-64/7.1/lib	osx86-64 –mcmodel=small libs
/opt/pgi/osx86-64/7.1/include	osx86-64 header files
/opt/pgi/osx86-64/7.1/man	UNIX-style man pages
/opt/pgi/osx86-64/7.1/src	Fortran 90 source files for included modules.

2.6.3 End-user Environment Settings on Mac OS

Now that you have installed the compilers in, for example, /opt/pgi, you must initialize your environment to use the compilers successfully. Each user must issue the following sequence of commands to initialize the shell environment before using the PGI compilers and tools.

The following commands make the 32-bit compilers the default.

- In bash, zsh, or ksh, use these commands:
 - \$ PGI=/opt/pgi; export PGI
 - \$ MANPATH=\$MANPATH:\$PGI/osx86/7.1/man
 - \$ export MANPATH
 - \$ PATH=\$PGI/osx86/7.1/bin:\$PATH
 - \$ export PATH
- In csh, use these commands:

```
% setenv PGI /opt/pgi
```

- % setenv MANPATH "\$MANPATH":\$PGI/osx86/7.1/man
- % set path = (\$PGI/osx86/7.1/bin \$path)

To make the 64-bit compilers the default, use these commands:

• In bash, zsh, or ksh, use these commands:

```
$ PGI=/opt/pgi; export PGI
$ MANPATH=$MANPATH:$PGI/osx86-64/7.1/man
$ export MANPATH
$ PATH=$PGI/osx86-64/7.1/bin:$PATH
$ export PATH
```

- In csh, use these commands:
 - % setenv PGI /opt/pgi
 - % setenv MANPATH "\$MANPATH":\$PGI/osx86-64/7.1/man
 - % set path = (\$PGI/osx86-64/7.1/bin \$path)

PGI Release 7.1 Release Notes

This document describes changes between previous releases and Release 7.1 of the PGI compilers, as well as late-breaking information not included in the current printing of the *PGI User's Guide*. There are nine platforms supported by the *PGI Workstation* and *PGI Server* compilers and tools:

- *32-bit Linux* supported on *32-bit Linux operating systems* running on either a 32-bit *x86* compatible or an *x64* compatible processor.
- 64-bit/32-bit Linux includes all features and capabilities of the 32-bit Linux version, and is also supported on 64-bit Linux operating systems running an x64 compatible processor.
- 32-bit Windows supported on 32-bit Windows operating systems running on either a 32-bit x86 compatible or an x64 compatible processor.
- 64-bit/32-bit Windows includes all features and capabilities of the 32-bit Windows version, and is also supported on 64-bit Windows operating systems running an x64 compatible processor.
- *32-bit SFU* supported on *32-bit Windows operating systems* running on either a 32-bit *x86* compatible or an *x64* compatible processor.
- *32-bit SUA* supported on *32-bit Windows operating systems* running on either a 32-bit *x86* compatible or an *x64* compatible processor.
- 64-bit/32-bit SUA includes all features and capabilities of the 32-bit SUA version, and is also supported on 64-bit Windows operating systems running an x64 compatible processor.
- *32-bit Apple Mac OS X* supported on 32-bit Apple Mac operating systems running on either a 32-bit or 64-bit Intel-based Mac system.

• *64-bit Apple Mac OS X* – supported on 64-bit Apple Mac operating systems running on a 64-bit Intel-based Mac system.

These release notes distinguish these versions where necessary.

3.1 PGI Release 7.1 Contents

Release 7.1 of PGI Workstation and PGI Server includes the following components:

- PGF95 native OpenMP and auto-parallelizing Fortran 95 compiler.
- PGF77 native OpenMP and auto-parallelizing FORTRAN 77 compiler.
- *PGHPF* data parallel High Performance Fortran compiler. Note: *PGHPF is not supported in Windows environments.*
- *PGCC* native OpenMP and auto-parallelizing ANSI C99 and K&R *C* compiler.
- *PGC*++ native OpenMP and auto-parallelizing ANSI *C*++ compiler.
- PGPROF multi-thread, OpenMP and MPI graphical profiler.
- *PGDBG* multi-thread, OpenMP and MPI graphical debugger.
 - *MPICH MPI libraries, version 1.2.7,* for both 32-bit and 64-bit development environments (Linux only)
 - Complete online documentation in PDF, HTML and UNIX man page formats.
 - A UNIX-like shell environment for *Win32* and *Win64* environments.

Depending on the product you purchased, you may not have licensed all of the above components.

3.2 Supported Systems

3.2.1 Supported Processors

The following table contains the processors on which Release 7.1 of the PGI compilers and tools is supported. The -tp < target > command-line option generates executables that utilize features and optimizations specific to a given CPU and operating system environment. Compilers included in a 64-bit/32-bit PGI installation can produce executables targeted to any 64-bit or 32-bit

target, including cross-targeting for AMD and Intel 64-bit AMD64 compatible CPUs.

In addition to the capability to generate binaries optimized for specific AMD or Intel processors, the PGI 7.1 compilers can produce PGI Unified Binary object or executable files containing code streams fully optimized and supported for both AMD and Intel x64 CPUs. The -tp < target > option must be used to produce unified binary files.

Processors Supported by PGI 7.1 Memory Floating Point HW Brand CPU Cores Address <target> x87 SSE1 SSE2 SSE3 64-bit Yes AMD Opteron/Quadcore 4 Yes Yes barcelona-64 Yes AMD barcelona 32-bit Yes Yes Yes Opteron/Quadcore 4 Yes 2 k8-64 AMD Opteron/Athlon64 32-bit Yes Yes Yes No AMD Opteron/Athlon64 2 k8-32 32-bit Yes Yes Yes No Opteron Rev E/F AMD Turion 2 k8-64e 64-bit Yes Yes Yes Yes /Athlon64 AMD Opteron Rev E/F 2 k8-32 32-bit Yes Yes Yes No Turion64 AMD Turion 1 k8-64e 64-bit Yes Yes Yes Yes /Athlon64 AMD Turion64 k8-32 32-bit Yes Yes Yes No 1 Core 2 2 32-bit Yes Yes Yes Intel core2 Yes Core 2 Intel 2 core2-64 64-bit Yes Yes Yes Yes Intel P4/Xeon EM64T 2 p7-64 64-bit Yes Yes Yes Yes Intel P4/Xeon EM64T 2 32-bit Yes Yes Yes Yes p7 Intel Xeon/Pentium4 1 32-bit Yes Yes Yes p7 No AMD Athlon XP/MP 1 32-bit Yes Yes No athlonxp No Pentium III Intel 1 32-bit Yes Yes No No piii AMD Athlon 1 athlon 32-bit Yes No No No AMD K6 1 32-bit Yes No No k6 No Intel Pentium II 1 p6 32-bit Yes No No No Generic x86 32-bit Yes No No Generic 1 No p5 or px

The table also includes the CPUs available and supported in multi-core versions.

3.2.2 Supported Operating Systems

The table lists the operating systems, and their equivalents, that Release 7.1 of the PGI compilers and tools supports.

To determine if Release 7.1 will install and run under a Linux equivalent version, such as Mandrake*, Debian*, Gentoo*, and so on, check the table for a supported system with the same glibc and gcc versions. Version differences in other operating system components can cause difficulties, but often these can be overcome with minor adjustments to the PGI software installation or operating system environment.

- Newer distributions of the Linux and Windows operating systems include support for x64 compatible processors and are designated 64-bit in the table. These are the only distributions on which the 64-bit versions of the PGI compilers and tools will fully install.
- If you attempt to install the 64-bit/32-bit Linux version on a system running a 32-bit Linux distribution, only the 32-bit PGI compilers and tools are installed.
- If you attempt to install the 64-bit Windows version on a system running 32-bit Windows, the installation fails.

Most newer Linux distributions support the *Native POSIX Threads Library* (NPTL), a new threads library that can be utilized in place of the *libpthreads* library available in earlier versions of Linux. Distributions that include NPTL are designated in the table. Parallel executables generated using the *OpenMP* and auto-parallelization features of the PGI compilers automatically make use of NPTL on distributions when it is available. In addition, the *PGDBG* debugger is capable of debugging executables built using either NPTL or earlier thread library implementations.

Multi-processor AMD Opteron processor-based servers use a *NUMA* (Non-Uniform Memory Access) architecture in which the memory latency from a given processor to a given portion of memory can vary. Newer Linux distributions, including SuSE 9/10 and SLES 9/10, include NUMA libraries that can be leveraged by a compiler and associated runtime libraries to optimize placement of data in memory.

In the table headings, HT = hyper-threading, NPTL = Native POSIX Threads Library, and NUMA = Non-Uniform Memory Access. For more information on these terms, see Terms and Definitions on page 2.

Operating Systems and Features Supported in PGI 7.1									
Distribution	Туре	64- bit	HT	pgC++	pgdbg	NPTL	NUMA	glibc	GCC
RHEL 5.0	Linux	Yes	Yes	Yes	Yes	Yes	No	2.6	4.1.2
RHEL 4.0	Linux	Yes	Yes	Yes	Yes	Yes	No	2.3.4	3.4.3
Fedora 7	Linux	Yes	Yes	Yes	Yes	Yes	Yes	2.6	4.1.2
Fedora 6	Linux	Yes	Yes	Yes	Yes	Yes	Yes	2.5	4.1.1
Fedora 5	Linux	Yes	Yes	Yes	Yes	Yes	Yes	2.4	4.1.0
Fedora 4	Linux	Yes	Yes	Yes	Yes	Yes	No	2.3.5	4.0.0
Fedora 3	Linux	Yes	Yes	Yes	Yes	Yes	No	2.3.3	3.4.2
Fedora 2	Linux	Yes	Yes	Yes	Yes	Yes	No	2.3.3	3.3.3
SuSE 10.3	Linux	Yes	Yes	Yes	Yes	Yes	Yes	2.6.1	4.2.1
SuSE 10.2	Linux	Yes	Yes	Yes	Yes	Yes	Yes	2.5	4.1.0
SuSE 10.1	Linux	Yes	Yes	Yes	Yes	Yes	Yes	2.4	4.1.0
SuSE 10.0	Linux	Yes	Yes	Yes	Yes	Yes	Yes	2.3.5	4.0.2
SuSE 9.3	Linux	Yes	Yes	Yes	Yes	Yes	Yes	2.3.4	3.3.5
SuSE 9.2	Linux	Yes	Yes	Yes	Yes	Yes	Yes	2.3.3	3.3.4
SLES 10	Linux	Yes	Yes	Yes	Yes	Yes	Yes	2.4	4.1.0
SLES 9	Linux	Yes	Yes	Yes	Yes	No	Yes	2.3.3	3.3.3
SuSE 9.1	Linux	Yes	Yes	Yes	Yes	Yes	No	2.3.3	3.3.3
RHEL 3.0	Linux	Yes	Yes	Yes	Yes	Yes	No	2.3.2	3.2.3
SuSE 9.0	Linux	Yes	Yes	Yes	Yes	No	No	2.3.2	3.3.1
SuSE 8.2	Linux	Yes	Yes	Yes	Yes	No	No	2.3.2	3.3
Red Hat 9.0	Linux	No	No	Yes	Yes	Yes	No	2.3.2	3.2.2
	ХР	No	Yes	Yes	Yes	NA	Yes	NA	NA
	2003	No	No	Yes	Yes	NA	Yes	NA	NA
Microsoft Windows	XP x64	Yes	Yes	Yes	Yes	NA	Yes	NA	NA
	2003 x64	Yes	Yes	Yes	Yes	NA	Yes	NA	NA
	SFU	No	Yes	Yes	Yes	NA	Yes	SFU	3.3
	SUA x86	No	Yes	Yes	Yes	NA	Yes	SUA	3.3
	SUA X04	Yes	Yes	Yes	Yes	NA NA	Yes	SUA NA	5.5 NA
	VISTA	res	res	res	res	INA	res	NA	INA
Apple	Core 2	No	No	Yes	Yes	NA	NA	NA	4.0.1
Mac OS X	Core 2 Duo	Yes	No	Yes	Yes	NA	NA	NA	4.0.1

Note. www.pgroup.com/support/install.htm lists any new Linux, Apple or Windows distributions that may be explicitly supported by the PGI compilers. If your operating system is newer than any of those listed in the preceding table, the installation may still be successful.

3.3 New or Modified Compiler Features

Following are the new features of Release 7.1 of the PGI compilers and tools as compared to prior releases.

- Multiple PGI Unified Binary Targets The -tp switch now supports a comma-separated list of targets allowing programs to be optimized for more than two 64-bit targets. Unified Binary directives and pragmas support the latest processors including Barcelona and Core 2 Duo.
- Performance improvements 5% to 10% performance improvements measured on several standard benchmarks, research community codes and ISV codes (see this link for recent benchmark testing: <u>http://www.pgroup.com/benchmark</u>.
- *OpenMP* Unlimited OpenMP thread counts are available in all PGI configurations. The number of threads is now unlicensed in the OpenMP run-time libraries, the debugger, and the profiler up to the hard limit of 64 threads.
- *PGC++ low-overhead exceptions implementation –* A new implementation of exception handling defers the cost of exception handling until an exception is thrown. For programs with many exception regions and few throws, this option may lead to improved run-time performance. The new implementation is available on Windows and newer Linux systems with libgcc_eh.
- Improved memory allocation on Windows A wrapper library for the Windows system memory allocation routines is available with the -Msmartalloc options. The wrapper library configures the small-block heap size to improve performance for programs that allocate lots of memory. The memory allocator also attempts to align large blocks of memory to avoid cache conflicts.
- Use huge pages for memory allocation –Msmartalloc has been enhanced to support large TLBs on Linux and Windows. For programs that access a lot of memory, huge pages may reduce the number of TLB misses and improve performance. This option is most effective on Barcelona and Core 2 systems; older architectures may have less benefit due to fewer available TLB entries.
- New Fortran intrinsics, subroutines, and attributes Implemented GET_COMMAND_ARGUMENT, GET_ENVIRONMENT_VARIABLES, GET_COMMAND, LEADZ, POPCNT, POPPAR, SIZEOF, CTIME8 and TIME8. On Windows, added GETDAT, GETTIM, TIME64, CTIME64, and LOCALTIME64 and support for the attribute, DECORATE.

- *Additional Fortran I/O features* Added ACCESS=APPEND parameter and support for the FLUSH statement.
- *OpenMP enhancements* Implemented C99/C++ style parallel *for* loops and added 64-bit long, long long, and all unsigned types as index variables.
- *Windows Environment* Expanded coverage of Windows library routines for inter-procedural analysis (IPA). Support long command lines and linking programs with many arguments with response files (@filename) on the command line.
- Enhanced gcc/g++ compatibility Support for the extension
 __builtin_expect and the attributes aligned, weak, and alias.
 Added support for variables in specified registers and controlling names
 used in assembler code.
- *Enhanced profile-feedback and code placement* optimizations Profilebased feedback optimizations and code layout is available on 64- and 32-bit platforms for F95, F77, C, and C++. New optimizations based on profile feedback include call-site inlining, switch-statement expansion, and register allocation. Static code layout is enabled on all platforms.
- *Extended SSE vectorization*—Loops with indirect addressing are now vectorized and FMAX, FMIN, DMAX and DMIN reductions are recognized. More loops with single-to-double precision conversions are vectorized.
- *Reduced use of temporary variables* Reduce the size of the stack frame by reusing temporary arrays and especially the associated descriptors when copying subprogram actual arguments. Eliminate copies when returning values from array-valued functions.
- *FLEXIm licensing upgrade* FLEXIm version 11.4.1 is included on all platforms.
- *Expanded OS support* Added support for Fedora Core 7 and SuSE 10.3 On Windows, for native compilation and execution under SFU and SUA using the Gnu *ld*.
- *Quad-Core AMD Opteron Processor support* –Updated tuning parameters and memory copy routines for Barcelona Rev B. Updated the opcode for Barcelona POPCNT from 'OF B8' (Rev A) to 'F3 OF B8' (Rev B) in the Windows assembler.
- *Apple Core 2 and Core 2 Duo support* PGI Workstation for *Mac OS X* is available from http://www.pgroup.com.

3.4 Compiler Options

3.4.1 Getting Started

By default, the PGI 7.1 compilers generate code that is optimized for the type of processor on which compilation is performed, the compilation host. If you are unfamiliar with the PGI compilers and tools, a good option to use by default is –fast or -fastse.

3.4.1.1 Using -fast, -fastsse, and Other Performance-Enhancing Options

These options create a generally optimal set of flags for targets that support SSE/SSE2 capability. These options incorporate optimization options to enable use of vector streaming SIMD (SSE/SSE2) instructions for 64-bit targets. They enable vectorization with SEE instructions, cache alignment, and flushz.

Note. The contents of the -fast and -fastsse options are host-dependent.

	<i>y y i i</i>
• -fast and -	fastsse typically include these options:
-02	Specifies a code optimization level of 2.
-Munroll=c:1	Unrolls loops, executing multiple instances of the loop during each iteration.
-Mnoframe	Indicates to not generate code to set up a stack frame
-Mlre.	Indicates loop-carried redundancy elimination
• These add - <i>fast</i> for 6	itional options are also typically available when using 4-bit targets and <i>-fastsse</i> for both 32- and 64-bit targets:
-Mvect=sse	Generates SSE instructions
-Mscalarsse	Generates scalar SSE code with xmm registers; implies -Mflushz
-Mcache_align	Aligns long objects on cache-line boundaries
-Mflushz	Sets SSE to flush-to-zero mode
-M[no]vect	Controls automatic vector pipelining.

Note. For best performance on processors that support SSE instructions, use the *PGF95* compiler, even for FORTRAN 77 code, and the *-fastsse* option.

In addition to *fast and fastsse*, the *Mipafast* option for interprocedural analysis and optimization can improve performance. You may also be able to obtain further performance improvements by experimenting with the individual *Mpgflag* options detailed in the *PGI User's Guide*, such as *Mvect*, *Munroll*, *Minline*, *Mconcur*, *Mpfi/Mpfo*, and so on. However, increased speeds using these options are typically application- and system-dependent. It is important to time your application carefully when using these options to ensure no performance degradations occur.

3.4.2 New or Modified Compiler Options

Unknown options are now treated as errors instead of warnings. This change makes it a compiler error to pass switches that are not known to the compiler; however, you can use the new switch *–noswitcherror* to issue warnings instead of errors for unknown switches.

The following compiler options have been added or modified in PGI 7.1:

- *-Bdynamic* Compiles for and links to the DLL version of the PGI runtime libraries. For more information, refer to Static and Dynamic Linking on page 64.
- *-Bstatic* Compiles for and links to the static version of the PGI runtime libraries. For more information, refer to Static and Dynamic Linking on page 64.
- -fast The -fast option adds –Mautoinline for C/C++ on 64-bit targets.
- -*Minfo=[no]file* The -*Minfo* option adds a sub-option, -*[no]file*, which prints source file names as they are compiled. The default is *Minfo=file*.
- -noswitcherror Issue warnings instead of errors for unknown switches. This behavior can be configured in the siterc file with set NOSWITCHERROR=1
- -*Minline=[no]reshape* Allow (disallow) inlining in Fortran even when array shapes do not match. The default is -*Minline=noreshape*, except with -*Mconcur* or -*mp*, when the default is -*Minline=reshape*.
- -*Mipa=[no]reshape* Enable (disable) Fortran inline with mismatched array shapes.

- -*Mmpi=mpich1* Use the default MPICH-1 communication libraries. Replaces the –Mmpi option, which is deprecated. Other MPI-related options are available with the PGI CDK.
- -*Mprof=mpich1* Use the profiled MPICH-1 coomincation library. Replaces the -Mprof=mpi option, which is deprecated.
- -M[no]smartalloc[=huge|huge:<n>|hugebss] -Msmartalloc has been enhanced to support large TLBs on Linux and Windows. This option must be used to compile the main routine to enable optimized malloc routines. The sub-option huge enables large 2 megabyte pages to be used by a single program. The effect is to reduce the number of TLB entries required to execute a program. This option is most effective on Barcelona and Core 2 systems; older architectures may have less benefit due to fewer available TLB entries.

The supported sub-options are:

hugeLink in the huge page runtime libraryhuge:<n>Limit the number of used pages to NhugebssPut the BSS section in huge pages

By itself, the huge sub-option will try to allocate as many huge pages as required. The number of huge pages can be limited with the :n sub-option or the environment variable PGI_HUGE_PAGES. The *hugebss* attempts to put a program's uninitialized data section into huge pages.

• -M[no]unroll[=c:<n>|n:<n>|m:<n>] - Added the capability to unroll loops with multiple blocks, particularly those with conditional statements, can now be unrolled. A new option to -Munroll, "m", is available to control this feature.

n: <n> Unroll single-block loops n times

m: <n> Unroll multi-block loops n times

The default setting is not to enable -Munroll = m. The default count with -Munroll = m is 4.

- -O4 The –O4 option now enables enhancements to algebraic transformations and register allocation.
- -*Mvect*[=[no]gather] The gather sub-option will vectorize loops containing indirect array references, such as:

```
sum = 0.d0
do k=d(j),d(j+1)-1
    sum = sum + a(k)*b(c(k))
enddo
```

The default is *-Mvect=gather*.

- -*M*[*no*]*propcond* Enable (disable) constant propagation from assertions derived from equality conditionals. Enabled by default.
- - [no]traceback Adds debug information for runtime traceback for use with the environment variable \$PGI_TERM. By default, traceback is enabled for f77 and f90 and disabled for C and C++. Setting set TRACEBACK=OFF; in a siterc or .mypg*rc will also disable default traceback. Using ON instead of OFF enabled default traceback.
- -*M*[*no*]*fpapprox*[=*div*|*sqrt*] Perform certain single-precision fp operations using low-precision approximation.

divApproximate floating point divisionsqrtApproximate floating point square rootrsqrtApproximate floating point reciprocal square rootBy default -Mfpapprox is not used. If -Mfpapprox is used without sub-options, it defaults to use approximate div, sqrt, and rsqrt.

- -M[no]fpmisalign Allow use of vector arithmetic instructions with memory operands whose addresses are not aligned on 16-byte boundaries for barcelona. On barcelona, an initialization routine to mask misalign fault on Barcelona processors guarded by runtime check of CPU ID. The default is -Mnofpmisalign on all processors, including barcelona. The option is effective if and only if it is used in conjunction with -tp barcelona-64 or -tp barcelona, or the code is compiled on a barcelona processor. The compiled code can only be run on a barcelona processor.
- -*M*[*no*]*loop32* Align innermost loops on 32-byte boundaries on barcelona. Small loops on barcelona may run fast if aligned on 32-byte boundaries; however, in practice, most assemblers do not yet implement efficient padding causing some programs to run more slowly with this option. Use -*Mloop32* on systems with an assembler tuned for barcelona. The default is -*Mnoloop32*.
- *-alias=[ansi|traditional]* Select optimizations based on type-based pointer alias rules in C and C++.

ansi Enable optimizations using ANSI C type-based pointer disambiguation

traditional Disable type-based pointer disambiguation For C, the default is *-alias=ansi*. For C++, the default is *alias=traditional*.

- -*Xs* Use legacy standard mode for C and C++; now also implies *alias=traditional*.
- -*Xt* Use legacy transitional mode for C and C++; now also implies *alias=traditional*.

- —*zc_eh* Generate zero-overhead exception regions. The —*zc_eh* option defers the cost of exception handling until an exception is thrown, so for program with many exception regions and few throws, this option may lead to improved run-time performance. The default is not to use *zc_eh* but instead to use —*sjlj_eh* which implements exception handling with *setjmp* and *longjmp*. This option is compatible with C++ code that was compiled with previous versions of PGI C++. The —*zc_eh* option is available only on newer Linux systems that supply the system unwind libraries in *libgcc_eh* and on Windows
- -*Xm* This option has been removed. It instructed the C++ compiler to allow dollar signs in names. Now, like always, dollar signs are accepted.
- -Bstatic and -Bdynamic On Windows, the -Bstatic and -Bdynamic options are new. Use -Bstatic to compile and link programs with static libraries. Use -Bdynamic to compile and link programs with DLLs. Note: the same option must be consistently used to compile and link all of the files that will be linked into an executable. The default is -Bstatic.
- -*Mmakedll* The -*Mmakedll* option now implies -*Mdynamic*.
- *-Mdll* The *-Mdll* option has been removed. Use *-Bdynamic* instead.
- -*stack=nocheck* The -stack option was changed to disable the automatic run-time stack extension feature on Windows. If the reserve and commit sub-options are set large enough to hold the full stack, no automatic extension is needed and the stack check can be avoided. The default is

-stack=check. On Win64, there is no default reserve or commit. On Win32, the reserve and commit default sizes are each 2,097,152 bytes.

- -Mchkstk Programs compiled with –Mchkstk are also instrumented to collect the stack high-water mark. If the environment variable PGI_STACK_USAGE is set at run time, the stack high-water mark will be printed at program exit.
- -[no]compress_names Compress C++ mangled names to fit into 1024 characters. Highly nested template parameters can cause very long function names. These long names can cause problems for older assemblers. The current default is -no_compress_names. All C++ user code must be recompiled when using this switch. Libraries supplied by PGI work with --compress_names.
- -*V* The -*V* option now prints the processor names. For example, on a Core 2 Duo, -*V* will print -tp core2-64.

• *-Mstandard* – *-Mstandard* now implies *-Mbackslash*, inhibiting recognizing the backslash escape sequences when -Mstandard is present, i.e., a backslash is treated as a normal character.

3.5 PGDBG New and Modified Features

PGI Workstation 7.1 includes several new features and enhancements in the *PGDBG* parallel debugger.

- PGDBG 7.1 now supports debugging of applications consisting of up to four MPICH-1 processes running on the same machine as PGDBG.
- PGDBG 7.1 supports debugging of MSMPI applications running on Microsoft Windows CCS clusters.
- PGDBG 7.1 enhancements including:
 - Improved stack trace capability
 - Improved interoperability with Microsoft Visual C++
 - Improved interoperability with gcc/g++
 - Fast disassembly and overall performance improvements

For a description of the usage and capabilities of PGDBG, see the *PGI Tools Guide*. For limitations and workarounds, see *www.pgroup.com/support/faq.htm*.

3.6 **PGPROF New and Modified Features**

PGI Workstation 7.1 includes several new features and enhancements in the *PGPROF* parallel profiler.

- PGPROF 7.1 now supports MPI profiling of collective communication routines.
- PGPROF 7.1 now supports MSMPI profiling on Microsoft Windows CCS clusters.
- PGPROF 7.1 includes some graphical user interface improvements.

For a description of the usage and capabilities of PGPROF, see the *PGI Tools Guide*. For limitations and workarounds, see *http://www.pgroup.com/support/fag.htm*.

3.7 Running an MPICH Program on Linux

PGI Workstation 7.1 for *Linux* includes MPICH libraries, tools, and licenses required to compile, execute, profile, and debug MPI programs. *PGI Workstation* can be installed on a single node, and the node can be treated as if it is a cluster. The MPI profiler and debugger are limited to processes on a single node in *PGI Workstation*. The *PGI CDK Cluster Development Kit* supports general development on clusters.

The PGI Tools Guide describes the MPI enabled tool in detail:

- PGPROF graphical MPI/OpenMP/multi-thread performance profiler.
- PGDBG graphical MPI/OpenMP/multi-thread symbolic debugger
- MPICH MPI libraries, version 1.2.7, for both 32-bit and 64-bit development environments (Linux only)

Try the MPI "hello world" program in the bench/mpihello subdirectory:

```
% cp -r $PGI/linux86/7.1/examples/mpi ./mpihello
% cd ./mpihello/mpihello
% pgf77 -o mpihello mpihello.f -Mmpi
% mpirun mpihello
Hello world! I'm node 0
% mpirun -np 4 mpihello
Hello world! I'm node 0
Hello world! I'm node 1
Hello world! I'm node 1
Hello world! I'm node 3
```

To run the debugger on an mpich program, do this:

```
mpirun -dbg=pgdbg -np 4 mpihello
```

3.8 Using the PGI Windows CDK with Microsoft Compute Cluster Server

If you have a PGI Windows CDK (Cluster Development Kit) license, then your PGI software includes support for working with Microsoft Compute Cluster Server and MSMPI. Specifically, this software includes support for these things:

• Building MPI applications with MSMPI

- Debugging cluster applications on a Windows CCS cluster with PGDBG
- Using PGPROF to do MPI profiling of MSMPI applications.

This section provides information on these tasks.

3.8.1 Build MPI Applications with MSMPI

Using MSMPI libraries

To build an application using the MSMPI libraries, use the -Mmpi=msmpi option. This option inserts options into the compile and link lines to pick up the MSMPI headers and libraries.

Note. The user can set the environment variables MPIDIR and MPILIBNAME to override the default values for the MPI directory and library name.

Generate MPI Profile Data

To build an application that generates MPI profile data, use the -Mprof=msmpi option. This option performs MPICH-style profiling for Microsoft MSMPI. For Microsoft Compute Cluster Server only, using this option implies -Mmpi=msmpi.

Note. For -Mprof=msmpi to work properly, the CCP_HOME environment variable must be set. This variable is typically set when the Microsoft Compute Cluster SDK is installed.

The profile data generated by running an application built with the option -Mmpi=msmpi contains information about the number of sends and receives, as well as the number of bytes sent and received, correlated with the source location associated with the sends and receives. -Mprof=msmpi must be used in conjunction with -Mprof=func or -Mprof=lines. When invoked using this type of profile data, PGPROF automatically displays MPI statistics.

3.8.2 Debug Cluster Applications that Generate MPI Profile Data

To invoke the PGDBG debugger to debug an MSMPI application, use the pgdbg <code>-mpi</code> option:

% pgdbg -mpi[:<path>] <mpiexec_args> [-program_args arg1,...argn]

The location of mpiexec should be part of your PATH environment variable. Otherwise, you should specify the pathname for mpiexec, or another similar launcher, as <path>in -mpi[:<path>].

To start a distributed debugging session, you must use the job submit command on the command line. You must also ensure that the debugger has access to the pgserv.exe remote debug agent on all nodes of the cluster used for the debug session.

To make pgserv.exe available, copy it from the PGI installation directory, such as C:\Program Files\PGI\win64\7.1-2\bin\pgserv.exe, into a directory or directories such that the path to pgserv.exe is the same on all nodes in the debug session.

Then you can start the debug session as follows:

% pgdbg -pgserv:<path_to_pgserv.exe> -mpi[:<job submit command>]

If you use a command similar to the following one, it assumes that a copy of pgserv.exe can be found in the current directory.

% pgdbg -pgserv -mpi[:<job submit command>]

Example

Suppose you wanted to debug the following job invocation:

```
"job.cmd" submit /numprocessors:4 /workdir:\\cce-
head\d\srt /stdout:sendrecv.out mpiexec sendrecv.exe
```

Use this command:

```
% pgdbg -pgserv "-mpi:c:\Program Files\Microsoft Compute
Cluster Pack\Bin\job.cmd" submit /numprocessors:4
/workdir:\\cce-head\d\srt /stdout:sendrecv.out mpiexec
sendrecv.exe
```

For this command to execute properly, a copy of pgserv.exe must appear in \\cce-head\d\srt.

Since the CCP installation updates the default PATH, the following command is equivalent to the previous one:

```
% pgdbg -pgserv -mpi:job.cmd submit /numprocessors:4
/workdir:\\cce-head\d\srt /stdout:sendrecv.out mpiexec
sendrecv.exe
```

Note. The use of quotes around the -mpi option varies, depending on the type of shell you are using. In the example, or if you are using cmd, specify the option as "-mpi:...", including the quotes around the option as well as around the optional job submit command. When invoking in a bash shell, you can specify the -mpi option as -mpi:"...", using the quotes around only the job submit command.

3.9 PGI Workstation 7.1 for Windows

PGI Workstation 7.1 for *Windows* supports most of the features of the 32- and 64-bit versions for *linux86* and *linux86-64* environments. The product optionally provides a familiar and somewhat compatible development environment for Linux or RISC/UNIX users porting to or developing programs for Windows systems. Except where noted in the *PGI User's Guide* and these release notes, the PGI compilers and tools on *Windows* function identically to their *Linux* counterparts.

PGI Workstation includes the *Microsoft Open Tools* tools, libraries, and include files. *Open Tools* includes C header files that use C++ style comments. Also, some Microsoft header files generate warnings about things such as multiple definitions of types. In most cases these warnings may be safely ignored.

3.9.1 The Windows Command Environment

A UNIX-like shell environment, *Cygwin*, is bundled with *PGI Workstation 7.1* for *Windows* to provide a familiar development environment for Linux or UNIX users. *PGI Workstation* for SFU and SUA does not include Cygwin; it runs in the SFU/SUA shell environment.

After installation, a double-left-click on the *PGI Workstation* icon on your desktop will launch a Cygwin bash shell command window with pre-initialized environment settings. Many familiar UNIX commands are available, such as vi, sed, grep, awk, make, and so on. If you are unfamiliar with the bash shell, refer to the user's guide included with the online HTML documentation.

On *Win64*, the desktop icon starts a bash shell configured for building 64-bit programs. To start a bash shell configured for building 32-bit programs, launch *PGI Workstation (32-bit)* from the Start menu.

Alternatively, by selecting the appropriate option from the *PGI Workstation* program group accessed in the usual way through the "Start" menu, you can launch a standard Windows command window that is pre-initialized to enable use of the PGI compilers and tools.

The command window launched by *PGI Workstation* can be customized using the "Properties" selection on the menu accessible by right-clicking the window's title bar.

3.9.2 MKS Toolkit Compatibility

The MKS Toolkit is a commercially available product providing a suite of UNIX and Windows utilities and is available for Win32 and Win64. You can use *PGI Workstation* compilers and tools in any of the MKS toolkit shells. To use *PGI Workstation* in an MKS shell, you must first configure your environment.

In the following example, the Windows system drive is C: and default installations were selected for the Java JRE and *PGI Workstation*. Open a Windows command prompt and execute the following commands:

```
> set PGI=C:\Program Files\PGI
> PATH=C:\Program Files (x86)\Java\j2re1.5.0_05/bin;%PATH%
> PATH=%PGI%\win64\7.1-2\bin;%PATH%
> set TMPDIR=C:\temp
```

Invoke an MKS shell. For example, to start the bash shell, type:

```
> bash.exe
```

For more information or to obtain the MKS Toolkit, visit the MKS website at http://www.mkssoftware.com/.

3.9.3 Using Shared object files in SFU and SUA

PGI Workstation for 32-bit SFU and 32-bit SUA now uses the GNU ld for its linker, unlike previous versions that used the Windows LINK.EXE. With this change, the PGI compilers and tools are now able to generate shared object (.so) files. You use the -shared switch to generate a shared object file.

The following example creates a shared object file, "hello.so", and then creates a program called "hello" that uses it.

Example 1:

First create a shared object file called "hello.so"

pgcc -shared hello.c -o hello.so

Then create a program that uses the shared object, in this example, "hello.so":

```
pgcc hi.c hello.so -o hello
```

When running a program that uses a shared object, you may encounter an error message similar to the following:

```
hello: error in loading shared libraries
hello.so: cannot open shared object file: No such file
or directory
```

This error message either means that the shared object file does not exist or its location is not specified in your LD_LIBRARY_PATH variable. To specify the location in your LD_LIBRARY_PATH variable, add the shared object's directory to your variable.

Example 2:

The following example adds the current directory to your LD_LIBRARY_PATH variable under C Shell, enter:

```
setenv LD_LIBRARY_PATH "$LD_LIBRARY_PATH":"./"
```

The following list is a summary of compiler switches that support shared objects:

-shared	Used to produce shared libraries
-Bdynamic	Passed to linker; specify dynamic binding
-Bstatic	Passed to linker; specify static binding
-Bstatic_pgi	Use to link static PGI libraries with dynamic system libraries; implies -Mnorpath
-L <libdir></libdir>	Passed to linker; Add directory to library search path
-Mnorpath	Don't add -rpath paths to link line
-Mnostartup	Do not use standard linker startup file
-Mnostdlib	Do not use standard linker libraries
-R <ldarg></ldarg>	Passed to linker; just link symbols from object, or add directory to run time search path

3.10 PGI Workstation 7.1 for Mac OS X

PGI Workstation 7.1 for *Mac OS X* supports most of the features of the 32and 64-bit versions for *linux86* and *linux86-64* environments. Except where noted in these release notes or the reference manuals, the PGI compilers and tools on *Mac OS X* function identically to their *Linux* counterparts.

3.10.1 Mac OS X Debugging Requirements

PGDBG must be a part of the procmod group in order to run. Either the user running the debugger must be a member of the procmod group or the debugger must be set-gid to the procmod group.

To set the debugger to be set-gid to the procmod group, execute the following commands:

```
su root
chgrp procmod $PGI/osx86/7.1-2/bin/pgdebug
chmod g+s $PGI/osx86/7.1-2/bin/pgdebug
```

Both the –g switch and the location of object files play important roles when compiling a program on Mac OS X for debugging.

- To debug a program with symbol information on the Mac OS X, files must be compiled with the -g switch to keep the program's object files, the files with a ".o" extension. Further, these object files must remain in the same directory in which they were created.
- If a program is built with separate compile and link steps, by compiling with the -c switch which generates the ".o" object files, then using the -g switch guarantees the required object files are available for debugging.

Use the following command sequence to compile and then link your code.

To compile the programs, use these commands:

```
pgcc -c -g main.c
pgcc -c -g foo.c
pgcc -c -g bar.c
```

To link, use this command:

pgcc -g main.o foo.o bar.o

3.11 Generating PGI Unified Binaries

All PGI compilers can produce PGI Unified Binary programs containing code streams fully optimized and supported for both AMD64 and Intel EM64T processors using the *-tp* target option. The compilers generate and combine into one executable multiple binary code streams each optimized for a specific platform. At runtime, this one executable senses the environment and dynamically selects the appropriate code stream.

Different processors have differences, some subtle, in hardware features such as instruction sets and cache size. The compilers make architecture-specific decisions about such things as instruction selection, instruction scheduling, and vectorization. PGI unified binaries provide a low-overhead means for a single program to run well on a number of hardware platforms.

Executable size is automatically controlled via unified binary culling. Only those functions and subroutines where the target affects the generated code will have unique binary images, resulting in a code-size savings of 10-90% compared to generating full copies of code for each target.

Programs can use PGI Unified Binary even if all of the object files and libraries are not compiled as unified binaries. Like any other object files, you can use PGI Unified Binary object files to create programs or libraries. No special start up code is needed; support is linked in from the PGI libraries.

The *-Mpfi* option disables generation of PGI Unified Binary. Instead, the default target auto-detect rules for the host are used to select the target processor.

3.11.1 Unified Binary Command-line Switches

The PGI Unified Binary command-line switch is an extension of the target processor switch, -tp, which may be applied to individual files during compilation.

The target processor switch, $-t_P$, accepts a comma-separated list of 64-bit targets and generates code optimized for each listed target.

The following example generates optimized code for three targets.

-tp k8-64,p7-64,core2-64

A special target switch, -tp x64, is the same as -tp k8-64, p7-64.

3.11.2 Unified Binary Directives and Pragmas

Unified binary directives and pragmas may be applied to functions, subroutines, or whole files. The directives and pragmas cause the compiler to generate PGI Unified Binary code optimized for one or more targets. No special command line options are needed for these pragmas and directives to take effect.

The syntax of the Fortran directive is

```
!pgi$[g|r| ] pgi tp [target]...
```

where the scope is g (global), r (routine) or blank. The default is r, routine.

For example,

!pgi\$g pgi tp k8_64 p7_64

indicates that the whole file, represented by g, should be optimized for both $k8_{64}$ and $p7_{64}$.

The syntax of the C/C++ pragma is

#pragma [global|routine|] tp [target]...

where the scope is global, routine, or blank. The default is routine.

For example, the following syntax indicates that the next function should be optimized for $k8_{64}$, $p7_{64}$, and $core2_{64}$.

```
#pragma routine tp k8_64 p7_64 core2_64
```

3.12 Static and Dynamic Linking on Windows

Prior to the 7.1 release, on Windows, all executables were linked against the PGI runtime DLL called pg.dll and the multi-threaded DLL version of the Microsoft runtime libraries. All executables were therefore dependent upon pg.dll and the Microsoft runtime. PGI now provides two compiler options that allow you to select static or dynamic linking.

3.12.1 -Bdynamic

Use this option to compile for and link to the DLL version of the PGI runtime libraries. This flag is required when linking with any DLL built by the PGI compilers. This flag corresponds to the /MD flag used by Microsoft's cl compilers.

On Windows, -Bdynamic must be used for both compiling and linking.

When you use the PGI compiler flag -Bdynamic to create an executable that links to the DLL form of the runtime, the executable built is smaller than one built without -Bdynamic. The PGI runtime DLLs, however, must be available on the system where the executable is run. The -Bdynamic flag must be used when an executable is linked against a DLL built by the PGI compilers.

3.12.2 -Bstatic

Use this option to explicitly compile for and link to the static version of the PGI runtime libraries.

On Windows, -Bstatic must be used for both compiling and linking.

For more information and usage examples for *–Bdynamic* and *–Bstatic* as well as information on using and creating static and dynamically-linked libraries, refer to the *PGI User's Guide*.

3.13 Using Environment Modules

On Linux, if you use the Environment Modules package (e.g., the module load command), PGI 7.1 includes a script to set up the appropriate module files.

Assuming your installation base directory is /opt/pgi, and your MODULEPATH environment variable is

/usr/local/Modules/modulefiles, execute this command:

```
/opt/pgi/linux86/7.1-2/etc/modulefiles/pgi.module.install \
    -all -install /usr/local/Modules/modulefiles
```

This command creates module files for all installed versions of the PGI compilers. You must have write permission to the modulefiles directory to enable the module commands:

```
module load pgi32/7.1
module load pgi64/7.1
module load pgi/7.1
```

where "pgi/7.1" uses the 32-bit compilers on a 32-bit system and uses 64-bit compilers on a 64-bit system.

To see what versions are available, use this command:

module avail pgi

The module load command sets or modifies the environment variables as follows:

PGI	the base installation directory
CC	full path to pgcc
FC	full path to pgf90
F90	full path to pgf90
F77	full path to pgf77
СРР	full path to pgCC
CXX	path to pgCC
C++	path to pgCC
РАТН	prepends the PGI compiler and tools bin directory
MANPATH	prepends the PGI man page directory
LD_LIBRARY_PATH	prepends the PGI library directory

PGI does not support the Environment Modules package. For more information about the package, go to: *http://modules.sourceforge.net*.

3.14 The REDIST Directories

Programs built with PGI compiler may depend on run-time library files. These library files must be distributed with such programs to enable them to execute on systems where the PGI compilers are not installed. There are PGI redistributable files for all platforms. On Windows, PGI also supplies Microsoft redistributable files.

3.14.1 PGI Redistributables

The PGI 7.1 release includes these directories:

- \$PGI/linux86/7.1/REDIST
- \$PGI/linux86-64/7.1/REDIST
- \$PGI/win64/7.1-2/REDIST
- \$PGI/win32/7.1/REDIST

These directories contain all of the PGI Linux runtime library shared object files or Windows dynamically linked libraries that can be re-distributed by PGI 7.1 licensees under the terms of the PGI End-user License Agreement (EULA). For reference, a copy of the PGI EULA is included in the 7.1 directory in text form.

The Linux REDIST directories contain the PGI runtime library shared objects for all supported targets. This enables users of the PGI compilers to create packages of executables and PGI runtime libraries that will execute successfully on almost any PGI-supported target system, subject to these requirements:

- End-users of the executable have properly initialized their environment
- On Linux, users have set LD_LIBRARY_PATH to use the relevant version of the PGI shared objects.

3.14.2 Microsoft Redistributables

The PGI products on Windows include Microsoft Open Tools. The Microsoft Open Tools directory contains a subdirectory named "redist". PGI 7.1 licensees may redistribute the files contained in this directory in accordance with the terms of the PGI End-User License Agreement.

3.15 Customizing With siterc and User rc Files

The PGI 7.1 release for Linux platforms includes a siterc file in the bin directory to enable site-specific customization of the PGI compiler drivers. Using siterc, you can control how the compiler drivers invoke the various components in the compilation tool chain. In addition to the siterc file, user *rc* files can reside in a given user's home directory: .mypgf77rc,

.mypgf90rc, .mypgccrc, .mypgcpprc, and .mypghpfrc can be used to control the respective PGI compilers. All of these files are optional.

The following table contains some examples that show how these rc files can be used to tailor a given installation for a particular purpose.

To perform this task	Do this:
Make the libraries found in /opt/newlibs/64 available to all linux86-64 compilations	Add the line: set SITELIB=/opt/newlibs/64; to /opt/pgi/linux86-64/7.1/bin/siterc
Make the libraries found in /opt/newlibs/32 available to all linux86 compilations.	Add the line: set SITELIB=/opt/newlibs/32; to /opt/pgi/linux86/7.1/bin/siterc
Add a new library path /opt/local/fast to all linux86-64 compilations.	Add the line: append SITELIB=/opt/local/fast; to /opt/pgi/linux86-64/7.1/bin/siterc
Make the include path /opt/acml/include available to all compilations; - <i>I/opt/acml/include</i> .	Add the line: set SITEINC=/opt/acml/include; to /opt/pgi/linux86/7.1/bin/siterc and to /opt/pgi/linux86-64/7.1/bin/siterc
Change – <i>Mmpi</i> to link in / <i>opt/mympi/64/libmpix.a</i> with <i>linux86-64</i> compilations.	Add the line: set MPILIBDIR=/opt/mympi/64; set MPILIBNAME=mpix; to /opt/pgi/linux86-64/7.1/bin/siterc
Have <i>linux86-64</i> compilations always add <i>–DIS64BIT –DAMD</i>	Add the line: set SITEDEF=IS64BIT AMD; to /opt/pgi/linux86-64/7.1/bin/siterc
To perform this task	Do this:
-------------------------------------	--
A user builds an F90	Add the line:
executable for <i>linux86-64</i> or	set RPATH=./REDIST;
<i>linux86</i> that resolves PGI	to ~/.mypgf95rc.
shared objects in the relative	NOTE: this will only affect the behavior
directory ./ <i>REDIST</i>	of PGF95 for the given user.

3.16 Known Limitations

The frequently asked questions (FAQ) section of the *pgroup.com* web page at *http://www.pgroup.com/support/index.htm* provides more up-to-date information about the state of the current release.

- When debugging a MPI job that is launched under pgserv, the processes in the job are stopped before the first instruction of the program. Since there is no source level debugging information at this point, issuing the source level next command executes very slowly. To avoid having to run the job until it completes, stops due to an exception, or stops by a PGDBG halt command entered by the user, the user should set an initial breakpoint. If a Fortran program is being debugged, set the initial breakpoint at main, or MAIN_, or at another point on the execution path before issuing the continue command.
- Object and module files created using *PGI Workstation 7.1* compilers are incompatible with object files from *PGI Workstation 5.x* and prior releases.
- Object files compiled with *–Mipa* using *PGI Workstation 6.1* and prior releases must be recompiled with *PGI Workstation 7.1*.
- On Windows, the version of *vi* included in cygwin can have problems when the SHELL variable is defined to something it does not expect. In this case, the following messages appear when vi is invoked:

```
E79: Cannot expand wildcards
E79: Cannot expand wildcards
E79: Cannot expand wildcards
Hit ENTER or type command to continue
```

To workaround this problem, set SHELL to refer to a shell in the cygwin bin directory, e.g. /bin/bash.

• The *-i8* option can make programs incompatible with MPI, use of any INTEGER*8 array size argument can cause failures with these libraries.

- The –i8 option can make programs incompatible with the bundled ACML library. Visit *developer.amd.com* to check for compatible libraries.
- Programs that incorporate object files compiled using -mcmodel=medium cannot be statically linked. This is a limitation of the *linux86-64* environment, not a limitation specific to the PGI compilers and tools.
- Using -*Mipa=vestigial* in combination with -*Mipa=libopt* with PGCC, you may encounter unresolved references at link time. This problem is due to the erroneous removal of functions by the *vestigial* sub-option to -*Mipa*. You can work around this problem by listing specific sub-options to -*Mipa*, not including *vestigial*.
- Using *-Mprof=func*, *-mcmodel=medium* and *-mp* together on any of the PGI compilers can result in segmentation faults by the generated executable. These options should not be used together.
- Programs compiled and linked for *gprof*-style performance profiling using *-pg* can result in segmentation faults on system running version 2.6.4 Linux kernels.
- *OpenMP* programs compiled using *-mp* and run on multiple processors of a *SuSE 9.0* system can run very slowly. These same executables deliver the expected performance and speed-up on similar hardware running *SuSE 9.1* and above.
- ACML 3.6 is built using the *-fastsse* compile/link option, which includes *-Mcache_align*. When linking with ACML using the *-lacml* option on 32-bit targets, all program units must be compiled with *-Mcache_align*, or an aggregate option such as *-fastsse* which incorporates *-Mcache_align*. This process is not an issue on 64-bit targets where the stack is 16-byte aligned by default. The lower-performance, but fully portable, *libblas.a* and *liblapack.a* libraries can be used on CPUs that do not support SSE instructions.
- When compiling with *-fPIC* and linking with *-lacml*, you may get the message "error while loading shared libraries: libacml_mv.so: cannot open shared object file: No such file or directory." In this case, you must add *-lacml_mv* library to the link line.

- Using -*Mpfi* and -*mp* together is not supported. The -*Mpfi* flag will disable -*mp* at compile time, which can cause run-time errors in programs that depend on interpretation of OpenMP directives or pragmas. Programs that do not depend on OpenMP processing for correctness can still use profile feedback. The -*Mpfo* flag does not disable OpenMP processing.
- On Windows, runtime libraries built for debugging (e.g. msvcrtd and libcmtd) are not included with *PGI Workstation*. When a program is linked with *-g*, for debugging, the standard non-debug versions of both the PGI runtime libraries and the Microsoft runtime libraries are always used. This limitation does not affect debugging of application code.
- Dynamic Link Libraries (DLLs) built on the Microsoft Windows platform by the *PGI Workstation 7.1* compilers have the following known limitations:
 - DLLs cannot be produced with the *PGI Workstation* C++ compiler.
 - If a DLL is built with the *PGI Workstation* compilers, the runtime DLLs must be used. The compiler option *–Mmakedll* ensures the correct runtime libraries are used.
 - If an executable is linked with any *PGI Workstation*-compiled DLL, the *PGI Workstation* runtime library DLLs must be used; this means the static libraries, which are used by default, cannot be used. To accomplish this, use the compiler option *–Bdynamic* when creating the executable.
- *PGPROF* Do not use –*Mprof* with *PGI Workstation* runtime library DLLs. To build an executable for profiling, use the static libraries. When the compiler option –*Bdynamic* is *not* used, the static libraries are the default.
- PGPROF Times reported for multi-threaded sample-based profiles, that is, profiling invoked with –pg or –Mprof=time options, are for the master thread only. PGI-style instrumentation profiling with –
 Mprof={lines | func} or hardware counter-based profiling using –
 Mprof=hwcts must be used to obtain profile data on individual threads.
- *PGDBG* The watch family of commands is unreliable when used with local variables. Calling a function or subroutine from within the scope of the watched local variable may cause missed events and/or false positive events. Local variables may be watched reliably if program scope does not leave the scope of the watched variable. Using the watch family of commands with global or static variables is reliable.

- *PGDBG* The call command does not support the following F90/F95 features: array-valued functions, pointer-valued functions, assumed-shape array arguments, or pointer arguments. There is no known workaround to this limitation.
- *PGDBG* Before PGDBG can set a breakpoint in code contained in a shared library, .so or .dll, the shared library must be loaded.
- *PGDBG* Debugging of unified binaries, that is, programs built with the *-tp=x64* option, is not fully supported. The names of some subprograms are modified in the creation of the unified binary, and PGDBG does not translate these names back to the names used in the application source code. For detailed information on how to debug a unified binary, see <u>http://www.pgroup.com/support/tools.htm</u>.
- *PGDBG Windows* In *PGDBG* on the *Windows* platform, use the forward slash ('/') character to delimit directory names in file path names.

Note. This requirement does not apply to the PGDBG debug command or to target executable names on the command line, although this convention will work with those commands.

- *PGDBG SFU/SUA* Due operating system limitations, PGDBG does not support hardware watchpoints, that is, the "hwatch" command, on SFU/SUA systems.
- *PGDBG SFU/SUA* Due to operating system limitations, PGDBG supports multi-thread debugging only with 32-bit SUA programs, with one restriction: once stopped, the process may be continued by continuing all threads, or a single thread, but not a partial set of the threads. Attempts to continue a partial set of threads results in the entire process, all threads, being continued.

3.17 Corrections

The following problems have been corrected in the *PGI Workstation 7.1* release. Most were reported in *PGI Workstation 7.0* or previous releases. Problems found in *PGI Workstation 7.0* may not have occurred in previous releases. A table is provided that describes the summary description of the problem. An *Internal Compiler Error* (ICE) is usually the result of checks the compiler components make on internal data structures, discovering inconsistencies that could lead to faulty code generation. For a complete and up-to-date list of TPRs fixed in recent releases of the PGI compilers and tools, see <u>www.pgroup.com/support/release_tprs.htm</u>.

3.17.1 Corrections in 7.1-2

The following TPRs have been resolved in 7.1-2.

TPR	Language / Tool	Description
4096	All	Want to be able to link to static or dynamic libraries on Windows
4123	32/64-bit Fortran	OpenMP error ' Statement not allowed in WORKSHARE construct'
4186	32/64-bit Fortran	Problem setting array bounds from intrinsic with constant strings
4196	32/64-bit Fortran	Parameters incorrectly initialized
4256	32/64-bit Fortran	Problem when array constructor depends on an outer loop variable
4315	64-bit C/C++	Boost needs access to GLIBC trunc and strptime
4317	32/64-bit Fortran	Compilation warnings emitted for valid code (PGF90-W-0164- Overlapping data initializations)
4331	32/64-bit Fortran	Array pointer assignment fails within OpenMP WORKSHARE
4332	32/64-bit Fortran	Mac OS X Block Data in an archive is not found at link time
4336	32/64-bit C/C++	Extended asm 'register' variables may lose stores

3.17.2 Corrections in 7.1-1

The following TPRs have been resolved in 7.1-1.

TPR	Language / Tool	Description
3195	32/64-bit Fortran	Would like function than translates Fortran 90 iostat error codes to error strings
3425	64-bit Fortran	Request for directives to ignore possible dependences in array assignments
3448	32/64-bit Fortran	Enhancement request - add support for LEADZ(),POPCNT(),POPPAR()
3564	32/64-bit Fortran	Request: make -Mbackslash the default, or make -Mstandard imply -Mbackslash
3595	32/64-bit C/C++	pgCC build seems to hang in IPA during link step

TPR	Language / Tool	Description
3602	32/64-bit C/C++	pgCC openmp program fails with schedule (static)
3717	32/64-bit C/C++	Enhancement request: support packed struct, attribute((packed)), for pgcc
3726	All	DWARF3: use DW_AT_MIPS_LINKAGE_NAME for mangled C++, Fortran 90 names
3754	64-bit C/C++	64-bit pgcc with -mp and omp parallel for cannot handle long iterators
3777	32/64-bit Fortran	Forward reference to module function from earlier module function in character length fails
3778	32/64-bit Fortran	Forward reference to module function from earlier module function in character length fails
3796	32/64-bit Fortran	Fortran use module, only:operator(.x.) gives spurious error message about .neqv.
3818	64-bit Fortran	Win64 -Mchkstk with -fPIC fails at link time
3928	32/64-bit Fortran	Fortran Dwarf2 problem - need DW_AT_type for function as dummy argument
3939	32/64-bit Fortran	Enhancement Request - User's example should not complain 'Non-constant expression'
3970	32/64-bit C/C++	C++ compiler fails (segmentation fault) with > 60 if/else in while loop
3982	32/64-bit Fortran	F90 with multiple directories with two .mod files of the same name confuses pgf90
4001	32/64-bit C/C++	pgcc #pragma omp for with 64-bit index variable not properly implemented
4046	32/64-bit Fortran	Multiple imports of pgf90 shared object using "dlopen" cause runtime failures.
4051	All	ScaLapack example needs revision to work with 6.2-5 CDK
4064	32/64-bit Fortran	Request for support of F2003 GET_COMMAND_ARGUMENT
4070	All	PDF documentation is not well formatted, margins too large
4079	32/64-bit Fortran	request for CVF 'sizeof' in fortran
4081	32-bit C/C++	C++ code compiled '-tp px -fastsee' causes internal compiler error "fr_decr_use: bad usect"

TPR	Language / Tool	Description
4086	64-bit C/C++	C++ compiler does not implement -Mfcon
4088	All	fstat64 fails for Win64 (Fortran)
4111	32-bit Fortran	On SUA32, -Mchkfpstk can give "fp stack is not empty" runtime errors
4128	32/64-bit Fortran	pgf90 fails to recognize constant PARAMETER array elements in array bounds
4132	32/64-bit C/C++	Feature request: define "PIC=1" when "-fPIC" is used.
4134	32/64-bit C/C++	Linking with shared and nonshared libs causes ' static object marked for destruction more than once'
4148	32/64-bit C/C++	pgcpp does not properly recognize "builtin_expect" (gcc compatibility)
4149	32-bit Fortran	Allow for non-English Administrator group name for Windows/SUA installations
4150	32/64-bit Fortran	Fortran Reads of integers in ascii files with (*) format should catch non-integer strings
4151	64-bit C/C++	pgcpp fails with internal compiler error on Win64 with -MD switch
4154	32/64-bit Fortran	time() and ctime() return integer*8 - not integer*4 as documentation indicates
4158	32/64-bit Fortran	Request for F2008 C_SIZEOF function
4160	32/64-bit Fortran	PGF90 fails to inline character function with spurious message about argument count
4163	32/64-bit Fortran	PGF90 PARAMETER array set by a reshaping another array fails
4167	32/64-bit Fortran	pgf90 does not support transfer function in parameter specification
4173	32/64-bit Fortran	OpenMP parallel region IF clause with(.false.) incorrectly runs in parallel
4175	All	Linking with -lacml or -lacml_mp requires -lacml_mv as well
4180	32/64-bit Fortran	pgf90 reporting error when attempting to inline pointer function
4185	32-bit C/C++	pgcc fails with internal compiler error with -mp/-Mconcur, 'Unable to allocate a register 22'

TPR	Language / Tool	Description
4191	32-bit Fortran	32-bit pgf77/pgf90 gives internal compiler error at -O2, 'Unable to allocate a register'
4200	32/64-bit Fortran	SAVE attribute assigned to allocatable, flagged as warning
4201	32/64-bit Fortran	SYSTEM_CLOCK does not have enough resolution
4202	32/64-bit Fortran	NINT() does not agree with other compiler output.
4203	32/64-bit Fortran	Fortran allocatable components of local derived type are not auto-deallocated
4205	32/64-bit Fortran	pghpf fails with internal compiler error - 'mk_mem_ptr_shape: extnt not subs'
4207	32/64-bit Fortran	Please refer to ISO/IEC 1539-1:1997 for Fortran reference
4208	All	64-bit 'Large-array' programming examples should add Win64 limitations
4213	64-bit C/C++	Compiler takes a long time to compile files with lots of functions
4214	32/64-bit Fortran	Fortran compiler fails with Internal compiler error in register allocation.
4219	32/64-bit Fortran	PGF90 USE,ONLY:OPERATOR(.x.) causes spurious error messages
4224	32/64-bit Fortran	OSX compilers do not support static linking
4226	32/64-bit C/C++	C+++ link causes execution error "static object has been marked for destruction more than once"
4233	32/64-bit Fortran	Allow expression using PARAMETER array element in dimension of another module array
4235	32/64-bit Fortran	Fortran RAN() function does not work properly with -r8
4241	32/64-bit C/C++	pgcc gives spurious error message for barrier in nested OpenMP parallel region
4244	32/64-bit Fortran	Nested OpenMP parallel regions should compile with -mp
4245	32-bit Fortran	pgf77/pgf90 fail with internal compiler error in register allocation, with -O -fPIC
4246	32/64-bit Fortran	Incorrect error message with use, only::usergeneric.
4247	32/64-bit C/C++	Add pgCC -a, like pgCC -A, but generating warnings, not errors

TPR	Language / Tool	Description
4248	32/64-bit C/C++	pgCC failure with -O3 -g undefined label in generated asm code
4249	32/64-bit Fortran	Fortran compiler should give error, not fail, with SUM(1:n)
4251	32/64-bit Fortran	There was no documentation for the \$PGI/target/version/src directory.
4253	32/64-bit Fortran	Compiler segfaults with Fortran PARAMETER initialized with implied DO
4254	32/64-bit Fortran	Fortran trim call and // in argument character length attribute causes runtime error
4257	32/64-bit Fortran	Allow non-default-kind logical argument to SUM MASK argument
4259	32/64-bit C/C++	Win32 compiler driver fails for large @objfilelist
4265	32/64-bit C/C++	pgCC -O2 -g example causes "Error: can't resolve `.text' .LBxxxx"
4266	32/64-bit Fortran	Host allocatable not auto-deallocated if only allocated in contained routine
4268	32/64-bit Fortran	Using MS\$ATTRIBUTES C without explicit interface
4271	32-bit Fortran	pgf77 with -fpic -tp px generated bad code accessing static variables
4274	All	The install script should modify the default value of \$PGI in lmgrd.rc to install directory.
4276	32-bit Fortran	Compiler fails with -mp with adjustable array in private clause of parallel region
4279	64-bit C/C++	asm() example with 'm' constraint was not restricted to memory
4280	32/64-bit C/C++	Bad assembly generated with C++ using -O -g
4283	32-bit Fortran	DEC\$ATTRIBUTE DLLIMPORT,ALIAS names should not be decorated (Windows)
4284	32-bit Fortran	Request to add GETDAT and GETTIM to lib3f
4286	64-bit Fortran	Fortran LOG function with -Ktrap=denorm generates floating point exception
4287	All	When upgrading a CDK installation with a new compiler release, installedk should not fail.

TPR	Language / Tool	Description
4288	All	User Guide should explain -Ktrap=inexact in more detail
4292	32/64-bit Fortran	Passing Dummy allocatable arrays not working
4298	32-bit Fortran	Failure when using -Munroll -tp px switches
4300	32/64-bit Fortran	Dwarf2 information was not being generated for Fortran module variables on Windows
4302	64-bit Fortran	pgf90 shape() fails with -Mlarge_arrays
4306	32/64-bit C/C++	pgcc -Mcpp -Bstatic causes the driver to call compiler with no input file
4308	32/64-bit Fortran	Fortran derived type dummy argument with initialized type component caused compiler failure
4310	32/64-bit C/C++	Add better messages when replacing loop by call to optimized runtime routine
4312	32-bit Fortran	Infinite loop in compiler with LRE example
4314	64-bit Fortran	Incorrect loop address used for Fortran dummy argument when expanded in vector loop
4316	64-bit Fortran	PGF90 syntax error issued for DO construct name immediately following USE
4318	64-bit Fortran	PGF95 fails with internal compiler error compiling program with error in deallocate statement
4319	64-bit Fortran	Use of -tp x64 -Mprof=func with contained subroutines yields link time undefined references

4

Contact Information and Documentation

You can contact The Portland Group at:

The Portland Group STMicroelectronics, Inc. Two Centerpointe Drive Lake Oswego, OR 97035 USA

The PGI User Forum is monitored by members of the PGI engineering and support teams as well as other PGI customers. The forum newsgroups may contain answers to commonly asked questions. Log in to the PGI website to access the forum:

http://www.pgroup.com/userforum/index.php

Or contact us electronically using any of the following means:

Fax: +1-503-682-2637 Sales: sales@pgroup.com Support: trs@pgroup.com WWW: http://www.pgroup.com

All technical support is by e-mail or submissions using an online form at *http://www.pgroup.com/support*. Phone support is not currently available.

Many questions and problems can be resolved at our frequently asked questions (FAQ) site at *http://www.pgroup.com/support/faq.htm*.

Online documentation is available at *http://www.pgroup.com/doc* or in your local copy of the documentation in the release directory *doc/index.htm*.