

PGI[®] COMPILERS & TOOLS

RELEASE NOTES FOR X86 CPUS AND TESLA GPUS

Version 2018




TABLE OF CONTENTS

Chapter 1. What's New in PGI 2018.....	1
1.1. What's New in 18.1.....	1
1.2. OpenMP.....	5
1.3. C++ Compiler.....	6
1.3.1. C++17.....	6
1.3.2. C++ and OpenACC.....	7
1.3.3. C++ Compatibility.....	7
1.4. Runtime Library Routines.....	7
1.5. Library Interfaces.....	8
1.6. Environment Modules.....	8
1.7. LLVM Code Generator.....	8
Chapter 2. Release Overview.....	10
2.1. Licensing.....	10
2.1.1. Licensing Terminology.....	10
2.1.2. Bundled License Key.....	11
2.1.3. Node-locked and Network Floating Licenses.....	11
2.2. Release Components.....	11
2.2.1. Additional Components.....	12
2.2.2. MPI Support.....	12
2.3. Terms and Definitions.....	12
2.4. Supported Platforms.....	13
2.5. Supported Operating System Updates.....	13
2.5.1. Linux.....	13
2.5.2. Apple macOS.....	13
2.5.3. Microsoft Windows.....	13
2.6. CUDA Toolkit Versions.....	14
2.7. Precompiled Open-Source Packages.....	15
2.8. Getting Started.....	16
Chapter 3. Distribution and Deployment.....	18
3.1. Application Deployment and Redistributables.....	18
3.1.1. PGI Redistributables.....	18
3.1.2. Linux Redistributables.....	18
3.1.3. Microsoft Redistributables.....	19
Chapter 4. Troubleshooting Tips and Known Limitations.....	20
4.1. Platform-specific Issues.....	20
4.1.1. Linux.....	20
4.1.2. Apple macOS.....	20
4.1.3. Microsoft Windows.....	20
4.2. Issues Related to Debugging.....	21
4.3. Profiler-related Issues.....	21

4.4. OpenACC Issues.....	21
Chapter 5. Contact Information.....	23

LIST OF TABLES

Table 1	Third-Party Software Security Updates for PGI version 18.1	5
Table 2	Typical -fast Options	17
Table 3	Additional -fast Options	17

Chapter 1.

WHAT'S NEW IN PGI 2018



Important The PGI 2018 Release includes updated FlexNet license management software to address a [security vulnerability](#). Users of any previous PGI release must update their FlexNet license daemons to enable PGI 18.1 and subsequent releases. See Third-Party Software Security Updates below and our [FlexNet Update FAQ](#) for more information.

Welcome to Release 2018 of the PGI compilers and tools!

If you read only one thing about this PGI release, make it this chapter. It covers all the new, changed, deprecated, or removed features in PGI products released this year. It is written with you, the user, in mind.

Every PGI release contains user-requested fixes and updates. We keep a complete list of these fixed [Technical Problem Reports](#) online for your reference.

1.1. What's New in 18.1

Key Features

Added support for Intel Skylake and AMD Zen processors, including support for the AVX-512 instruction set on the latest Intel Xeon processors.

Added full support for OpenACC 2.6.

Enhanced support for OpenMP 4.5 for multicore CPUs, including SIMD directives as tuning hints. OpenMP 4.5 is supported on Linux/x86 with the LLVM code generator (see below).

Added support for the CUDA 9.1 toolkit, including on the latest NVIDIA Volta V100 GPUs.

OpenACC and CUDA Fortran

Changed the default CUDA Toolkit used by the compilers to CUDA Toolkit 8.0.

Changed the default compute capability chosen by the compilers to cc35,cc50,cc60 .

Added support for CUDA Toolkit 9.1.

Added full support for the OpenACC 2.6 specification including:

- ▶ serial construct
- ▶ if and if_present clauses on host_data construct
- ▶ no_create clause on the compute and data constructs
- ▶ attach clause on compute, data, and enter data directives
- ▶ detach clause on exit data directives
- ▶ Fortran optional arguments
- ▶ acc_get_property, acc_attach, and acc_detach routines
- ▶ profiler interface

Added support for asterisk (*) syntax to CUDA Fortran launch configuration. Providing an asterisk as the first execution configuration parameter leaves the compiler free to calculate the number of thread blocks in the launch configuration.

Added two new CUDA Fortran interfaces, `cudaOccupancyMaxActiveBlocksPerMultiprocessor` and `cudaOccupancyMaxActiveBlocksPerMultiprocessorWithFlags`. These provide hooks into the CUDA Runtime for manually obtaining the maximum number of thread blocks which can be used in grid-synchronous launches, same as provided by the asterisk syntax above.

OpenMP

Changed the default initial value of `OMP_MAX_ACTIVE_LEVELS` from 1 to 16.

Added support for the **taskloop** construct's **firstprivate** and **lastprivate** clauses.

Added support for the OpenMP Performance Tools (OMPT) interface. Available with the LLVM code generator compilers on Linux.

C++

Added support for GNU interoperability through GCC 7.2.

Added partial support for C++17 including `constexpr if`, fold expressions, structured bindings, and several other C++17 features. See [C++17](#) for a complete list of supported features.

Fortran

Changed how the PGI compiler runtime handles Fortran array descriptor initialization; this change means any program using Fortran 2003 should be recompiled with PGI 18.1.

Libraries

Reorganized the Fortran cuBLAS and cuSolver modules to allow use of the two together in any Fortran program unit. As a result of this reorganization, any codes which use cuBLAS or cuSolver modules must be recompiled to be compatible with this release.

Added a new PGI math library, libpgm. Moved math routines from libpgc, libpgftnrtl, and libpgf90rtl to libpgm. This change should be transparent unless you have been explicitly adding libpgc, libpgftnrtl, or libpgf90rtl to your link line.

Added new fastmath routines for single precision scalar/vector sin/cos/tan for AVX2 and AVX512F processors.

Added support for C99 scalar complex intrinsic functions.

Added support for vector complex intrinsic functions.

Added environment variables to control runtime behavior of intrinsic functions:

MTH_I_ARCH={em64t,sse4,avx,avxfma4,avx2,avx512knl,avx512}

Override the architecture/platform determined at runtime.

MTH_I_STATS=1

Provide basic runtime statistics (number of calls, number of elements, percentage of total) of elemental functions.

MTH_I_STATS=2

Provide detailed call count by element size (single/double-precision scalar, single/double-precision vector size).

MTH_I_FAST={relaxed,precise}

Override compile time selection of fast intrinsics (the default) and replace with either the relaxed or precise versions.

MTH_I_RELAXED={fast,precise}

Override compile time selection of relaxed intrinsics (the default with `-Mfprelaxed=intrinsic`) and replace with either the fast or precise versions.

MTH_I_PRECISE={fast,relaxed}

Override compile time selection of precise intrinsics (the default with `-Kieee`) and replace with either the fast or relaxed versions.

Profiler

Improved the CPU Details View to include the breakdown of time spent per thread.

Added an option to let one select the PC sampling frequency.

Enhanced the NVLink topology to include the NVLink version.

Enhanced profiling data to include correlation ID when exporting in CSV format.

Operating Systems and Processors

Added support for the AMD Zen (EPYC, Ryzen) processor architecture. Use the `-tp=zen` compiler option to target AMD Zen explicitly.

Added support for the Intel Skylake processor architecture. Use the `-tp=skylake` compiler option to target Intel Skylake explicitly.

Added support for the Intel Knights Landing processor architecture. Use the `-tp=kn1` compiler option to target Intel Knights Landing explicitly.

LLVM Code Generator

Released a production version of the PGI Linux/x86-64 compilers with an LLVM code generator and OpenMP runtime; these compilers can be found in the PGI installation directory `linux86-64-llvm` installed alongside the default PGI x86-64 compilers in `linux86-64`. See [LLVM Code Generator](#) for more information.

License Management

Updated FlexNet Publisher license management software to v11.14.1.3. This update addresses several issues including:

- ▶ A security vulnerability on Windows. See Third-Party Software Security Updates below and the [FlexNet Update FAQ](#) for more information.
- ▶ Seat-count stability improvements on network floating license servers when borrowing licenses (`lmborrow`) for off-line use. For early return of borrowed seats, users should invoke the new `"-bv"` option for `lmborrow`. See our [license borrowing FAQ](#) for more information.



Important Users with PGI 2017 (17.x) or older need to update their license daemons to support 18.1 or newer. The new license daemons are backward-compatible with older PGI releases.

Deprecations and Eliminations

Dropped support for the following versions of Linux:

- ▶ CentOS 5 through 6.3
- ▶ Fedora 6 through 13
- ▶ openSUSE 11 through 13.1 and openSUSE Leap through 42.2
- ▶ RHEL through 6.3
- ▶ SLES 11
- ▶ Ubuntu 12.04

Dropped support for versions of glibc older than 2.12.

Dropped support for macOS version 10.9 (Mavericks).

Stopped including components from CUDA Toolkit version 7.5 in the PGI packages. CUDA 7.5 can still be targeted if one directs the compiler to a valid installation location of CUDA 7.5 using `CUDA_HOME`.

Deprecated legacy PGI accelerator directives. When the compiler detects a deprecated PGI accelerator directive, it will print a warning. This warning will include the OpenACC directive corresponding to the deprecated directive if one exists. Warnings about deprecated directives can be suppressed using the new `legacy` sub-option to the `-acc` compiler option. The following library routines have been deprecated: `acc_set_device`, `acc_get_device`, and `acc_async_wait`; they have been replaced by `acc_set_device_type`, `acc_get_device_type`, and `acc_wait`, respectively. The following environment variables have been deprecated: `ACC_NOTIFY` and `ACC_DEVICE`; they have been replaced by `PGI_ACC_NOTIFY` and `PGI_ACC_DEVICE_TYPE`, respectively. Support for legacy PGI accelerator directives may be removed in a future release.

Dropped support for CUDA x86. The `-Mcuda=x86` compiler option is no longer supported.

Dropped support for CUDA Fortran emulation mode. The `-Mcuda=emu` compiler option is no longer supported.

Third-Party Software Security Updates

Table 1 Third-Party Software Security Updates for PGI version 18.1

CVE ID	Description
CVE-2016-10395	Updated FlexNet Publisher to v11.14.1.3 to address a vulnerability on Windows. We recommend all users update their license daemons –see the FlexNet Update FAQ . For more information, see the Flexera website .

1.2. OpenMP

OpenMP 3.1

The PGI Fortran, C, and C++ compilers support OpenMP 3.1 on all platforms.

OpenMP 4.5

The PGI Fortran, C, and C++ compilers compile most OpenMP 4.5 programs for parallel execution across all the cores of a multicore CPU or server. **target** regions are implemented with default support for the multicore host as the target, and **parallel** and **distribute** loops are parallelized across all OpenMP threads. This feature is supported on Linux/x86 platforms with the LLVM code generator only.

Current limitations include:

- ▶ The **simd** construct can be used to provide tuning hints; the **simd** construct's **private**, **lastprivate**, **reduction**, and **collapse** clauses are processed and supported.
- ▶ The **declare simd** construct is ignored.
- ▶ The **ordered** construct's **simd** clause is ignored.
- ▶ The **task** construct's **depend** and **priority** clauses are not supported.
- ▶ The loop construct's **linear**, **schedule**, and **ordered(n)** clauses are not supported.
- ▶ The **declare reduction** directive is not supported.

1.3. C++ Compiler

1.3.1. C++17

The PGI 18.1 C++ compiler introduces partial support for the C++17 language standard; access this support by compiling with `--c++17` or `-std=c++17`.

Supported C++17 core language features are available on Linux (requires GCC 5 or later) and OS X.

This PGI compiler release supports the following C++17 language features:

- ▶ Structured bindings
- ▶ Selection statements with initializers
- ▶ Compile-time conditional statements, a.k.a. **constexpr if**
- ▶ Fold expressions
- ▶ Inline variables
- ▶ C++17 lambdas
- ▶ Lambda capture of `*this` by value

The following C++17 language features are not supported in this release:

- ▶ Class template deduction
- ▶ Auto non-type template parameters
- ▶ Guaranteed copy elision

The PGI products do not include a C++ standard library, so support for C++17 additions to the standard library depends on the C++ library provided on your system. On Linux, GCC 7 is the first GCC release with significant C++17 support. On OS X, there is no support for any of the C++17 library changes with one exception: `std::string_view` is available on OS X High Sierra.

The following C++ library changes are supported when building against GCC 7:

- ▶ `std::string_view`
- ▶ `std::optional`

- ▶ `std::variant`
- ▶ `std::any`
- ▶ Variable templates for metafunctions

The following C++ library changes are not available on any system that this PGI release supports:

- ▶ Parallel algorithms
- ▶ Filesystem support
- ▶ Polymorphic allocators and memory resources

1.3.2. C++ and OpenACC

There are limitations to the data that can appear in OpenACC data constructs and compute regions:

- ▶ Variable-length arrays are not supported in OpenACC data clauses; VLAs are not part of the C++ standard.
- ▶ Variables of class type that require constructors and destructors do not behave properly when they appear in data clauses.
- ▶ Exceptions are not handled in compute regions.
- ▶ Member variables are not fully supported in the `use_device` clause of a `host_data` construct; this placement may result in an error at runtime.

1.3.3. C++ Compatibility

Optional packages that provide a C++ interface, such as the MPI package included with all PGI products, require the use of the `pgc++` compiler driver for compilation and linking.

These optional packages include:

- ▶ ESMF
- ▶ MPICH
- ▶ MVAPICH
- ▶ NetCDF
- ▶ Open MPI
- ▶ Parallel NetCDF
- ▶ ScaLAPACK

1.4. Runtime Library Routines

PGI 2018 supports runtime library routines associated with the PGI Accelerator compilers. For more information, refer to *Using an Accelerator* in the PGI Compiler User's Guide.

1.5. Library Interfaces

The PGI products contain a number of libraries that export C interfaces by using Fortran modules. These libraries and functions are described in the PGI Compiler User's Guide.

1.6. Environment Modules

On Linux, if you use the Environment Modules package (e.g., the `module load` command), then PGI 2018 includes a script to set up the appropriate module files.

1.7. LLVM Code Generator

PGI 2018 includes a production version of the PGI Linux/x86-64 compilers with an LLVM code generator and OpenMP runtime. It includes most features, languages, and programming model support found in the default PGI compilers, as well as the PGI debugger and profiler.

You have a few options for using the PGI compilers with the LLVM code generator including the following:

Compiler Option: `-Mllvm`

If you set your path to include the location of the default compilers, `linux86-64/18.1/bin`, then the compilation option `-Mllvm` will invoke the LLVM back end compilers in `linux86-64-llvm/18.1/bin`.

Environment Modules: `pgi-llvm`

If your environment is set up to use PGI's environment module files, the following commands load the PGI compilers with the default code generator:

```
module load pgi/18.1
```

By contrast, to load the PGI compilers with the LLVM code generator, use this set of commands:

```
module load pgi-llvm
module load pgi/18.1
```

To clear your environment, you can use:

```
module purge
```

Environment variable: `PATH`

Set your path to include `linux86-64-llvm/18.1/bin`.

Take care not to mix object files compiled with the default PGI compilers and the PGI compilers with the LLVM backend. While the generated code is compatible, the OpenMP runtime libraries are not.

Features available in the PGI Linux/x86-64 compilers using the LLVM code generator:

- ▶ Optimized OpenMP atomics.
- ▶ OpenMP 4.5 features, not including GPU offload.

- ▶ Improved performance for some applications compared to the default PGI x86-64 code generator.

Limitations of the PGI Linux/x86-64 compilers with the LLVM code generator:

- ▶ Only available for Linux; not available for macOS or Windows.
- ▶ Fortran debugging is available only for breakpoints, call stacks, and basic types.
- ▶ PGI Unified Binary is not available.
- ▶ Interprocedural optimization using the `-Mipa` option is not available.
- ▶ Some source-code directives, e.g. `DEC$`, may not have an effect.

Chapter 2.

RELEASE OVERVIEW

This chapter provides an overview of Release 2018 of the PGI Accelerator™ C11, C++14 and Fortran 2003 compilers and development tools for 64-bit x86-compatible processor-based workstations, servers, and clusters running versions of the Linux, Apple macOS and Microsoft Windows operating systems.

2.1. Licensing

All PGI products include exactly the same PGI compilers and tools software. The difference is in which features are enabled by the license keys.

PGI release 2018 version 18.1 and newer contain updated (v11.14.1.3) [FlexNet Publisher](#) license management software.

This FlexNet update addresses a [security vulnerability](#) on Windows, [license borrowing](#) ('Imborrow') issues and other improvements.



Important Users with PGI 2017 (17.x) or older need to update their license daemons to support 18.1 or newer. The new license daemons are backward-compatible with older PGI releases. For more information, see the [FlexNet Update FAQ](#).

2.1.1. Licensing Terminology

The PGI compilers and tools are license-managed. Before discussing licensing, it is useful to have common terminology.

- ▶ **License** – the right to use PGI compilers and tools as defined by the End-user License Agreement (EULA), this is a legal agreement between NVIDIA and PGI end-users. PGI Professional (for-fee, perpetual) licenses are identified by a Product Identification Number (PIN - see below). You can find a copy of the EULA on the [PGI website](#), [pgicompile.com/LICENSE](#), and in the `$PGI/<platform>/<rel_number>/doc` directory of every PGI software installation.
- ▶ **License keys** – ASCII text strings that enable use of the PGI software and are intended to enforce the terms of the License. For PGI Professional, License keys are generated by each PGI end-user on the PGI website using a unique hostid and are

typically stored in a file called `license.dat` that is accessible to the systems for which the PGI software is licensed.

- ▶ **PIN** – Product Identification Number, a unique 6-digit number associated with a PGI Professional license. This PIN is included in your order confirmation. The PIN can also be found in your license key file after **VENDOR_STRING=**.
- ▶ **PIN tie code** – A unique 16-digit number associated with each license (PIN) that allows others to "tie" that license to their [PGI user account](https://pgicompilers.com/), pgicompilers.com/ account for administrative purposes. PGI Professional licensees can use their PIN tie code to share license administration capabilities with others in their organization.

2.1.2. Bundled License Key

Installation may place a temporary license key file named `license.dat` in the PGI installation directory if no such file already exists.

If you use a separate license server, for example

`LM_LICENSE_FILE=port@server.domain.com`, that supports this version, it is recommended that you remove or rename the license key file in the installation directory.

2.1.3. Node-locked and Network Floating Licenses

- ▶ Node-locked single-user licenses allow one user at a time to compile solely on the system on which both the PGI compilers and tools, and PGI license server are installed.
- ▶ Network floating licenses allow one or more users to use the PGI compilers and tools concurrently on any compatible client systems networked to a license server, that is, the system on which the PGI network floating license key(s) are installed. There can be multiple installations of the PGI compilers and tools on client systems connected to the license server; and client systems can use the license concurrently up to the maximum number of seats licensed for the license server.

2.2. Release Components

Release 2018 includes the following components:

- ▶ PGFORTRAN™ native OpenMP and OpenACC Fortran 2003 compiler.
- ▶ PGCC® native OpenMP and OpenACC ANSI C11 and K&R C compiler.
- ▶ PGC++® native OpenMP and OpenACC ANSI C++14 compiler.
- ▶ PGI Profiler® OpenACC, CUDA, OpenMP, and multi-thread graphical profiler.
- ▶ PGI Debugger® MPI, OpenMP, and multi-thread graphical debugger.
- ▶ Open MPI version 2.1.2 for 64-bit Linux including support for NVIDIA GPUDirect. Note that 64-bit linuxx86-64 MPI messages are limited to < 2 GB size each. As NVIDIA GPUDirect depends on InfiniBand support, Open MPI is also configured to use InfiniBand hardware if it is available on the system. InfiniBand support requires OFED 3.18 or later.
- ▶ MPICH libraries, version 3.2, for 64-bit macOS development environments.

- ▶ ScaLAPACK 2.0.2 linear algebra math library for distributed-memory systems for use with Open MPI, MPICH or MVAPICH, and the PGI compilers on 64-bit Linux and macOS for Intel 64 or AMD64 CPU-based installations.
- ▶ Microsoft HPC Pack 2012 MS-MPI Redistributable Pack (version 4.1) for 64-bit development environments (Windows only).
- ▶ BLAS and LAPACK library based on the customized OpenBLAS project source.
- ▶ A UNIX-like shell environment for 64-bit Windows platforms.
- ▶ FlexNet license utilities.
- ▶ Documentation in man page format and [online](https://pgicompilers.com/docs), pgicompilers.com/docs, in both HTML and PDF formats.

2.2.1. Additional Components

PGI floating license holders may download additional components for Linux from the PGI website including:

- ▶ MPICH MPI libraries
- ▶ MVAPICH2 MPI libraries

2.2.2. MPI Support

You can use PGI products to develop and debug MPI applications. PGI node-locked licenses support debugging of up to 16 local MPI processes. PGI network floating licenses provide the ability to debug up to 256 local or remote MPI processes.

2.3. Terms and Definitions

This document contains a number of terms and definitions with which you may or may not be familiar. If you encounter an unfamiliar term in these notes, please refer to the [PGI online glossary](https://pgicompilers.com/definitions) located at pgicompilers.com/definitions.

These two terms are used throughout the documentation to reflect groups of processors:

Intel 64

A 64-bit Intel Architecture processor with Extended Memory 64-bit Technology extensions designed to be binary compatible with AMD64 processors. This includes Intel Pentium 4, Intel Xeon, Intel Core 2, Intel Core 2 Duo (Penryn), Intel Core (i3, i5, i7), both first generation (Nehalem) and second generation (Sandy Bridge) processors, as well as Ivy Bridge, Haswell, Broadwell, and Skylake processors.

AMD64

A 64-bit processor from AMD[™] incorporating features such as additional registers and 64-bit addressing support for improved performance and greatly increased memory range. This term includes the AMD Athlon64[™], AMD Opteron[™], AMD Turion[™], AMD Barcelona, AMD Shanghai, AMD Istanbul, AMD Bulldozer, AMD Piledriver, and AMD Zen processors.

2.4. Supported Platforms

There are three platforms supported by the PGI compilers and tools for x86-64 processor-based systems.

- ▶ **64-bit Linux** – supported on 64-bit Linux operating systems running on a 64-bit x86 compatible processor.
- ▶ **64-bit macOS** – supported on 64-bit Apple macOS operating systems running on a 64-bit Intel processor-based Macintosh computer.
- ▶ **64-bit Windows** – supported on 64-bit Microsoft Windows operating systems running on a 64-bit x86-compatible processor.

2.5. Supported Operating System Updates

This section describes updates and changes to PGI 2018 that are specific to Linux, macOS, and Windows.

2.5.1. Linux

- ▶ CentOS 6.4 through 7.4
- ▶ Fedora 14 through 27
- ▶ openSUSE 13.2 through openSUSE Leap 42.3
- ▶ RHEL 6.4 through 7.4
- ▶ SLES 12 through SLES 12 SP 3
- ▶ Ubuntu 14.04, 16.04, 17.04, 17.10

2.5.2. Apple macOS

PGI 2018 for macOS supports most of the features of the version for Linux environments. Except where noted in these release notes or the user manuals, the PGI compilers and tools on macOS function identically to their Linux counterparts.

- ▶ The compilers, debugger, and profiler are supported on macOS versions 10.10.5 (Yosemite) through 10.13 (High Sierra).

2.5.3. Microsoft Windows

PGI products for Windows support most of the features of the PGI products for Linux environments. PGI products on all Windows systems include the Microsoft Open Tools but also require that a Microsoft Windows Software Development Kit (SDK) be installed prior to installing the compilers.



PGI 2018 requires the Windows 10 SDK, even on Windows 7, 8 and 8.1.

These Windows operating systems are supported in PGI 2018:

- ▶ Windows Server 2008 R2
- ▶ Windows 7
- ▶ Windows 8
- ▶ Windows 8.1
- ▶ Windows 10
- ▶ Windows Server 2012
- ▶ Windows Server 2016

2.6. CUDA Toolkit Versions

The PGI compilers use NVIDIA's CUDA Toolkit when building programs for execution on an NVIDIA GPU. Every PGI installation packages puts the required CUDA Toolkit components into a PGI installation directory called `2018/cuda`.

An NVIDIA CUDA driver must be installed on a system with a GPU before you can run a program compiled for the GPU on that system. PGI products do not contain CUDA Drivers. You must download and install the appropriate [CUDA Driver from NVIDIA](#). The CUDA Driver version must be at least as new as the version of the CUDA Toolkit with which you compiled your code.

The PGI tool `pgaccelinfo` prints the driver version as its first line of output. Use it if you are unsure which version of the CUDA Driver is installed on your system.

PGI 18.1 contains the following versions of the CUDA Toolkits:

- ▶ CUDA 8.0 (default)
- ▶ CUDA 9.0
- ▶ CUDA 9.1

By default, the PGI compilers in this release use the CUDA 8.0 Toolkit from the PGI installation directory. You can compile with a different version of the CUDA Toolkit using one of the following methods:

- ▶ Use a compiler option. The `cudaX.Y` sub-option to `-Mcuda` or `-ta=tesla` where `X.Y` denotes the CUDA version. For example, to compile a C file with the CUDA 9.1 Toolkit you would use:

```
pgcc -ta=tesla:cuda9.1
```

Using a compiler option changes the CUDA Toolkit version for one invocation of the compiler.

- ▶ Use an rcfile variable. Add a line defining `DEFCUDAVERSION` to the `siterc` file in the installation `bin/` directory or to a file named `.mypgirc` in your home directory. For example, to specify the CUDA 9.1 Toolkit as the default, add the following line to one of these files:

```
set DEFCUDAVERSION=9.1;
```

Using an rcfile variable changes the CUDA Toolkit version for all invocations of the compilers reading the rcfile.

By default, the PGI compilers use the CUDA Toolkit components installed with the PGI compilers and in fact most users do not need to use any other CUDA Toolkit installation

than those provided with PGI. Developers working with pre-release CUDA software may occasionally need to test with a CUDA Toolkit version not included in a PGI release. Conversely, some developers might find a need to compile with a CUDA Toolkit older than the oldest CUDA Toolkit installed with a PGI release. For these users, PGI compilers can interoperate with components from a CUDA Toolkit installed outside of the PGI installation directories.

PGI tests extensively using the co-installed versions of the CUDA Toolkits and fully supports their use. Use of CUDA Toolkit components not included with a PGI install is done with your understanding that functionality differences may exist.

To use a CUDA toolkit that is not installed with a PGI release, such as CUDA 7.5 with PGI 18.1, there are three options:

- ▶ Use the rcfile variable `DEFAULT_CUDA_HOME` to override the base default

```
set DEFAULT_CUDA_HOME = /opt/cuda-7.5;
```

- ▶ Set the environment variable `CUDA_HOME`

```
export CUDA_HOME=/opt/cuda-7.5
```

- ▶ Use the compiler compilation line assignment `CUDA_HOME=`

```
pgfortran CUDA_HOME=/opt/cuda-7.5
```

The PGI compilers use the following order of precedence when determining which version of the CUDA Toolkit to use.

- ▶ In the absence of any other specification, the CUDA Toolkit located in the PGI installation directory `2018/cuda` will be used.
- ▶ The rcfile variable `DEFAULT_CUDA_HOME` will override the base default.
- ▶ The environment variable `CUDA_HOME` will override all of the above defaults.
- ▶ A user-specified `cudaX.Y` sub-option to `-Mcuda` and `-ta=tesla` will override all of the above defaults and the CUDA Toolkit located in the PGI installation directory `2018/cuda` will be used.
- ▶ The compiler compilation line assignment `CUDA_HOME=` will override all of the above defaults (including the `cudaX.Y` sub-option).
- ▶ The environment variable `PGI_CUDA_HOME` overrides all of the above; reserve `PGI_CUDA_HOME` for advanced use.

2.7. Precompiled Open-Source Packages

Many open-source software packages have been ported for use with PGI compilers on Linux x86-64.

The following PGI-compiled open-source software packages are included in the PGI Linux x86-64 download package:

- ▶ OpenBLAS 0.2.19 – customized BLAS and LAPACK libraries based on the OpenBLAS project source.
- ▶ Open MPI 2.1.2 – open-source MPI implementation.
- ▶ ScaLAPACK 2.0.2 – a library of high-performance linear algebra routines for parallel distributed memory machines. ScaLAPACK uses Open MPI 2.1.2.

The following list of open-source software packages have been precompiled for execution on Linux x86-64 targets using the PGI compilers and are available to download from the [PGI website](https://pgi.compilers.com/downloads) at pgi.compilers.com/downloads.

- ▶ MPICH 3.2 – open-source MPI implementation.
- ▶ MVAPICH2 2.2 – open-source MPI implementation.
- ▶ ESMF 7.0.2 for Open MPI 2.1.2 – The Earth System Modeling Framework for building climate, numerical weather prediction, data assimilation, and other Earth science software applications.
- ▶ ESMF 7.0.2 for MPICH 3.2.
- ▶ ESMF 7.0.2 for MVAPICH2 2.2.
- ▶ NetCDF 4.5.0 for C++11 – A set of software libraries and self-describing, machine-independent data formats that support the creation, access, and sharing of array-oriented scientific data, written in C. Included in this package are the following components:
 - ▶ NetCDF-C++ 4.3.0 – C++ interfaces to NetCDF libraries.
 - ▶ NetCDF-Fortran 4.4.4 – Fortran interfaces to NetCDF libraries.
 - ▶ HDF5 1.10.1 – data model, library, and file format for storing and managing data.
 - ▶ CURL 7.46.0 – tool and a library (usable from many languages) for client-side URL transfers.
 - ▶ SZIP 2.1.1 – extended-Rice lossless compression algorithm.
 - ▶ ZLIB 1.2.11 – file compression library.
- ▶ NetCDF 4.4.1.1 for C++98 – includes all the components listed in NetCDF for C++11 above.
- ▶ Parallel NetCDF 1.9.0 for MPICH 3.2
- ▶ Parallel NetCDF 1.9.0 for MVAPICH2 2.2
- ▶ Parallel NetCDF 1.9.0 for Open MPI 2.1.2

In addition, these software packages have also been ported to PGI on Linux x86-64 but due to licensing restrictions, they are not available in binary format directly from PGI. You can find instructions for building them in the [Porting & Tuning Guides](#) section of the PGI website at pgi.compilers.com/tips.

- ▶ FFTW 2.1.5 – version 2 of the Fast Fourier Transform library, includes MPI bindings built with Open MPI 2.1.2.
- ▶ FFTW 3.3.7 – version 3 of the Fast Fourier Transform library, includes MPI bindings built with Open MPI 2.1.2.

For additional information about building these and other packages, please see the [Porting & Tuning Guides](#) section of the PGI website at pgi.compilers.com/tips.

2.8. Getting Started

By default, the PGI 2018 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 the aggregate option `-fast`.

Aggregate options incorporate a generally optimal set of flags for targets that support SSE capability. These options incorporate optimization options to enable use of vector streaming SIMD instructions. They enable vectorization with SSE instructions, cache alignment, and flushz.



The content of the `-fast` option is host-dependent.

The following table shows the typical `-fast` options.

Table 2 Typical `-fast` Options

Use this option...	To do this...
<code>-O2</code>	Specifies a code optimization level of 2.
<code>-Munroll=c:1</code>	Unrolls loops, executing multiple instances of the original loop during each iteration.
<code>-Mnoframe</code>	Indicates to not generate code to set up a stack frame. Note With this option, a stack trace does not work.
<code>-Mlre</code>	Indicates loop-carried redundancy elimination.
<code>-Mpre</code>	Indicates partial redundancy elimination

`-fast` also typically includes the options shown in the following table:

Table 3 Additional `-fast` Options

Use this option...	To do this...
<code>-Mvect=simd</code>	Generates packed SSE and AVX instructions.
<code>-Mcache_align</code>	Aligns long objects on cache-line boundaries.
<code>-Mflushz</code>	Sets flush-to-zero mode.



For best performance on processors that support SSE and AVX instructions, use the PGFORTRAN compiler, even for FORTRAN 77 code, and the `-fast` option.

In addition to `-fast`, the `-Mipa=fast` 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 that are described in the *PGI Compiler Reference Manual*, 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.

Chapter 3.

DISTRIBUTION AND DEPLOYMENT

Once you have successfully built, debugged and tuned your application, you may want to distribute it to users who need to run it on a variety of systems. This section addresses how to effectively distribute applications built using PGI compilers and tools.

3.1. Application Deployment and Redistributables

Programs built with PGI compilers may depend on runtime 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 Linux and Windows. On Windows, PGI also supplies Microsoft redistributable files.

3.1.1. PGI Redistributables

The PGI 2018 Release includes these directories:

```
$PGI/linux86-64/18.1/REDIST
$PGI/win64/18.1/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 2018 licensees under the terms of the PGI End-User License Agreement (EULA). For reference, a text-form copy of the PGI EULA is included in the 18.1 `doc` directory.

3.1.2. Linux Redistributables

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.
- ▶ Users have set `LD_LIBRARY_PATH` to use the relevant version of the PGI shared objects.

3.1.3. Microsoft Redistributables

The PGI products on Windows include Microsoft Open Tools. The Microsoft Open Tools directory contains a subdirectory named `redist`. PGI 2018 licensees may redistribute the files contained in this directory in accordance with the terms of the PGI End-User License Agreement.

Microsoft supplies installation packages, `vcredist_x86.exe` and `vcredist_x64.exe`, containing these runtime files. These files are available in the `redist` directory.

Chapter 4.

TROUBLESHOOTING TIPS AND KNOWN LIMITATIONS

This section contains information about known limitations, documentation errors, and corrections. Wherever possible, a work-around is provided.

For up-to-date information about the state of the current release, please see the [PGI frequently asked questions \(FAQ\)](#) webpage.

4.1. Platform-specific Issues

4.1.1. Linux

The following are known issues on Linux:

- ▶ Programs that incorporate object files compiled using `-mmodel=medium` cannot be statically linked. This is a limitation of the linux86-64 environment, not a limitation of the PGI compilers and tools.

4.1.2. Apple macOS

The following are known issues on Apple macOS:

- ▶ The PGI 2018 compilers do not support static linking of binaries. For compatibility with future Apple updates, the compilers only support dynamic linking of binaries.

4.1.3. Microsoft Windows

The following are known issues on Windows:

- ▶ For the Cygwin `emacs` editor to function properly, you must set the environment variable **CYGWIN** to the value "tty" before invoking the shell in which `emacs` will run. However, this setting is incompatible with the PGBDG command line interface (`-text`), so you are not able to use `pgdbg -text` in shells using this setting.

- ▶ 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 Hit ENTER or type command to continue

To work around this problem, set **SHELL** to refer to a shell in the Cygwin `bin` directory, e.g., `/bin/bash`.

- ▶ On Windows, runtime libraries built for debugging (e.g., `msvcrtd` and `libcmtd`) are not included with PGI products. 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.

4.2. Issues Related to Debugging

The following are known issues in the PGI debugger:

- ▶ Debugging of PGI Unified Binaries, that is, programs built with more than one `-tp` option, is not fully supported. The names of some subprograms are modified in compilation and the debugger does not translate these names back to the names used in the application source code.
- ▶ When debugging on the Windows platform, the Windows operating system times out **stepi/nexti** operations when single stepping over blocked system calls.

4.3. Profiler-related Issues

Some specific issues related to the PGI Profiler:

- ▶ The Profiler relies on being able to directly call `'dlsym'`. If this system call is intercepted by the program being profiled or by some other library the profiler may hang at startup. We have encountered this specific problem with some implementations of MPI. We recommend you disable any features that may be intercepting the `'dlsym'` system call or disable CPU profiling with the `--cpu-profiling off` option.
 - ▶ To disable `'dlsym'` interception when using IBM's spectrum MPI omit the following option: `-gpu` and add the options `-x PAMI_DISABLE_CUDA_HOOK=1` and `-disbale_gpu_hooks`.

4.4. OpenACC Issues

This section includes known limitations in PGI's support for OpenACC directives. PGI plans to support these features in a future release.

ACC routine directive limitations

- ▶ Fortran assumed-shape arguments are not yet supported.

Clause Support Limitations

- ▶ Not all clauses are supported after the `device_type` clause.

Chapter 5.

CONTACT INFORMATION

You can contact PGI at:

20400 NW Amberwood Drive Suite 100
Beaverton, OR 97006

Or electronically using any of the following means:

Fax: +1-503-682-2637

Sales: [mailto: sales@pgroup.com](mailto:sales@pgroup.com)

WWW: <https://www.pgroup.com> or pgicompilers.com

The [PGI User Forum](#), pgicompilers.com/userforum is monitored by members of the PGI engineering and support teams as well as other PGI customers. The forums contain answers to many commonly asked questions. [Log in to the PGI website](#), pgicompilers.com/login to access the forums.

Many questions and problems can be resolved by following instructions and the information available in the [PGI frequently asked questions \(FAQ\)](#), pgicompilers.com/faq.

Submit support requests using the [PGI Technical Support Request](#) form, pgicompilers.com/support-request.

Notice

ALL NVIDIA DESIGN SPECIFICATIONS, REFERENCE BOARDS, FILES, DRAWINGS, DIAGNOSTICS, LISTS, AND OTHER DOCUMENTS (TOGETHER AND SEPARATELY, "MATERIALS") ARE BEING PROVIDED "AS IS." NVIDIA MAKES NO WARRANTIES, EXPRESSED, IMPLIED, STATUTORY, OR OTHERWISE WITH RESPECT TO THE MATERIALS, AND EXPRESSLY DISCLAIMS ALL IMPLIED WARRANTIES OF NONINFRINGEMENT, MERCHANTABILITY, AND FITNESS FOR A PARTICULAR PURPOSE.

Information furnished is believed to be accurate and reliable. However, NVIDIA Corporation assumes no responsibility for the consequences of use of such information or for any infringement of patents or other rights of third parties that may result from its use. No license is granted by implication of otherwise under any patent rights of NVIDIA Corporation. Specifications mentioned in this publication are subject to change without notice. This publication supersedes and replaces all other information previously supplied. NVIDIA Corporation products are not authorized as critical components in life support devices or systems without express written approval of NVIDIA Corporation.

Trademarks

NVIDIA, the NVIDIA logo, Cluster Development Kit, PGC++, PGCC, PGDBG, PGF77, PGF90, PGF95, PGFORTRAN, PGHPF, PGI, PGI Accelerator, PGI CDK, PGI Server, PGI Unified Binary, PGI Visual Fortran, PGI Workstation, PGPROF, PGROUP, PVF, and The Portland Group are trademarks and/or registered trademarks of NVIDIA Corporation in the U.S. and other countries. Other company and product names may be trademarks of the respective companies with which they are associated.

Copyright

© 2013-2018 NVIDIA Corporation. All rights reserved.