

PGI[®] COMPILERS & TOOLS

INSTALLATION AND RELEASE NOTES FOR
OPENPOWER CPUS

Version 2017



TABLE OF CONTENTS

Chapter 1. Release Overview.....	1
1.1. About This Release.....	1
1.2. Release Components.....	1
1.3. Supported Platforms.....	2
1.4. Supported Operating Systems.....	2
1.5. Supported CUDA Software.....	2
1.5.1. Targeting a CUDA Toolkit Version.....	2
1.6. Precompiled Open-Source Packages.....	2
1.7. Getting Started.....	3
Chapter 2. Compiler Features.....	5
2.1. What's New in Release 2017.....	5
2.2. Programming Models.....	7
2.3. Command-line Environment.....	7
2.4. Fortran Language.....	7
2.5. C Language.....	7
2.6. C++ Language.....	8
2.7. OpenACC.....	8
Chapter 3. Installation and Configuration.....	10
3.1. License Management.....	10
3.2. Environment Initialization.....	11
3.3. Network Installations.....	11
Chapter 4. Troubleshooting Tips and Known Limitations.....	13
4.1. Release Specific Limitations.....	13
Chapter 5. Contact Information.....	14

LIST OF TABLES

Table 1 Typical -fast Options4

Chapter 1.

RELEASE OVERVIEW

Welcome to Release 2017 of the PGI Accelerator™ C11, C++14 and Fortran 2003 compilers hosted on and targeting OpenPOWER+Tesla processor-based servers and clusters running versions of the Linux operating system.

1.1. About This Release

These PGI compilers generate host CPU code for 64-bit little-endian OpenPOWER CPUs, and GPU device code for NVIDIA Kepler and Pascal GPUs.

These compilers include all GPU OpenACC features available in the PGI C/C++/Fortran compilers for x86-64.

Documentation includes the `pgcc`, `pgc++` and `pgfortran` man pages and the `-help` option. In addition, you can find online versions of these installation and release notes, the *PGI Compiler User's Guide* and *PGI Compiler Reference Manual* for OpenPOWER. Online documentation is available at <http://www.pgroup.com/resources/docs.htm>.

1.2. Release Components

Release 2017 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 profiler.
- ▶ Open MPI version 1.10.2 including support for NVIDIA GPUDirect. GPUDirect requires CUDA 7.5 or later. 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.
- ▶ ScaLAPACK 2.0.2 linear algebra math library for distributed-memory systems for use with Open MPI and the PGI compilers.
- ▶ BLAS and LAPACK library based on the customized OpenBLAS project source.
- ▶ Documentation in man page format and online PDFs.

1.3. Supported Platforms

These OpenPOWER hardware/software platforms have been used in testing:

- ▶ CPUs: POWER8, POWER8E, POWER8NVL
- ▶ Linux distributions:
 - ▶ Ubuntu 14.04, 16.04
 - ▶ RHEL 7.3
- ▶ GCC versions: 4.8.4, 4.8.5, 5.3.1
- ▶ CUDA Toolkit versions:
 - ▶ 7.5 driver versions 352.39, 352.79
 - ▶ 8.0 driver version 361.93.02

1.4. Supported Operating Systems

The PGI 17.4 compilers were built on an OpenPOWER system running Ubuntu 14.04 OS with a GCC 4.8.2 toolchain. They have been tested on that platform, Ubuntu 16.04 with GCC 5.3.1, and RHEL 7.3 with GCC 4.8.5.

1.5. Supported CUDA Software

Select components of the CUDA Toolkit 7.5 and 8 are included under the PGI installation tree in `/opt/pgi`.

1.5.1. Targeting a CUDA Toolkit Version

- ▶ The CUDA 7.5 Toolkit is set as the default in PGI 17.4. To use the CUDA 7.5 Toolkit, first download the CUDA 7.5 driver from NVIDIA at <http://www.nvidia.com/cuda>.
- ▶ You can compile with the CUDA 8.0 Toolkit either by adding the option `-ta=tesla:cuda8.0` to the command line or by adding `set DEFCUDAVERSION=8.0` to the `siterc` file.
- ▶ `pgaccelinfo` prints the driver version as the first line of output. For a 7.5 driver, it prints:
CUDA Driver Version 7050

1.6. Precompiled Open-Source Packages

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

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

- ▶ OpenBLAS 0.2.19 – customized BLAS and LAPACK libraries based on the OpenBLAS project source.
- ▶ Open MPI 1.10.2 – 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 1.10.2.

The following list of open-source software packages have been precompiled for execution on OpenPOWER targets using the PGI compilers and are available to download from the [PGI website](#):

- ▶ MPICH 3.2 – open-source MPI implementation.
- ▶ NetCDF 4.4.1.1 – 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 1.10.0-patch1 – 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 – extended-Rice lossless compression algorithm.
 - ▶ ZLIB 1.2.8 – file compression library.
- ▶ Parallel NetCDF 1.7.0 for MPICH
- ▶ Parallel NetCDF 1.7.0 for Open MPI

In addition, these software packages have also been ported to PGI on OpenPOWER 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.

- ▶ FFTW 2.1.5 – version 2 of the Fast Fourier Transform library, includes MPI bindings built with Open MPI 1.10.2.
- ▶ FFTW 3.3.4 – version 3 of the Fast Fourier Transform library, includes MPI bindings built with Open MPI 1.10.2.

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

1.7. Getting Started

By default, the PGI 2017 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 that enable use of SIMD instructions .



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

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

Table 1 Typical `-fast` Options

Use this option...	To do this...
<code>-O2</code>	Specifies a code optimization level of 2 and <code>-Mvect=SIMD</code> .
<code>-Munroll=c:1</code>	Unrolls loops, executing multiple instances of the original loop during each iteration.
<code>-Mlre</code>	Indicates loop-carried redundancy elimination.
<code>-Mautoinline</code>	Enables automatic function inlining in C & C++.



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

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`, 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 2.

COMPILER FEATURES

Many user-requested fixes and updates are implemented in each PGI release. Refer to http://www.pgroup.com/support/release_tprs.htm for a complete and up-to-date table of technical problem reports fixed in recent releases of PGI compilers and tools. This table contains a summary description of each problem as well as the version in which it was fixed.

2.1. What's New in Release 2017

17.4 Updates and Additions

- ▶ PGI Accelerator Compilers
 - ▶ Added support for atomic add and atomic subtract in device code for Fortran's single precision complex data types.
 - ▶ Added device versions of `isfinite`, `isnan`, `isinf`, and `round`.
 - ▶ Added support for CUDA Fortran warp-vote operations.
 - ▶ Changed the impact of `-g` on the optimization level used when compiling device code. For host code, `-g` sets the compiler's optimization level to zero unless another optimization level is specified using a `-O` option. For device code, the same is now true. For details about generating debug information for device code, refer to the PGI Compiler Reference Guide's section on DWARF Debugging Formats.
 - ▶ Enabled `-ta=tesla:lineinfo`.
 - ▶ Updated the version of `nvlink` to 8.0.73; this version of the device linker allows an arbitrary ordering of object files on its command line.

17.3 Updates and Additions

Added the PGI C/C++ compiler plug-in for Eclipse versions Mars and Luna.

17.1 Updates and Additions

- ▶ Improved inlining with the `-Minline` option. You may see different functions inlined with this release than you observed in prior releases. In some cases, compilation may take longer to complete because of an increase in the number of functions inlined. Some of the `-Minline` sub-options, which you can use to adjust which functions get inlined, have also changed from previous releases:
 - ▶ Added `totalsize:n` to limit inlining to the total size of `n` where `n` is the size of the combined program units on a per file basis.
 - ▶ Changed `size:n` to `maxsize:n` which allows inlining only of functions smaller than approximately `n` lines. The compilers silently convert the previous `size:n` to `maxsize:n`.
 - ▶ Dropped `levels:n` which limited inlining to `n` levels of functions. The compilers silently ignore `levels:n`.
- ▶ Added the `-cpp` option as an alias for the `-Mpreprocess` option.
- ▶ Improved exception handling support and dropped the following related compiler options: `--[no_]exceptions`, `--need_exception_spec`, `--sjlj_eh`, `--[no]zc_eh`.
- ▶ Added the `-M[no]variadic_macros` to allow or disallow variadic macros. Variadic macros are allowed by default.
- ▶ Changed the default version of LLVM to 3.9.
- ▶ Improved support for `-ta=multicore`.
- ▶ Dropped support for CUDA 7.0.
- ▶ Changed the default version of the CUDA Toolkit used by the compilers from CUDA 7.0 to 7.5. The CUDA Toolkit 8.0 can be used instead of 7.5 by adding the sub-option `cuda8.0` to the `-ta=tesla` or `-Mcuda compile-` and link-time options.
- ▶ Improved support for the OpenACC **cache** directive.
- ▶ Added support for additional OpenACC 2.5 features:
 - ▶ Changed the behavior of the **exit data** directive to decrement the dynamic reference count.
 - ▶ Added the new optional **finalize** clause to set the dynamic reference count to zero.
 - ▶ Added the **if_present** clause to the **update** directive which changes the behavior when data is not present from a runtime error to a no-op.
 - ▶ Added new **init**, **shutdown**, and **set** directives.
 - ▶ Added new API routines to get and set the default async queue value.
 - ▶ Added support for the new definition of **routine bind** clause.
 - ▶ Updated the value of `_OPENACC` to 201510.

With the exception of nested parallelism, declare link, and adding restrictions to **cache** clause variable refs to variables within cached region, OpenACC 2.5 feature support is complete in PGI 2017.

2.2. Programming Models

CUDA Fortran and OpenACC in Fortran, C and C++ targeting Tesla are fully supported in this Release.

OpenMP 3.1 is fully supported in Fortran, C and C++. OpenMP 4.5 directives are processed by the Fortran compiler and all data scoping is supported, but many of the new directives have no other effect on code generation. In particular, there is no support for offloading using target directives. All target regions are scheduled for execution on the multicore OpenPOWER host processor. OpenMP 4.5 pragmas are not yet processed or supported in the C and C++ compilers.

2.3. Command-line Environment

The PGI compilers for OpenPOWER are command-line compatible with the corresponding PGI products on Linux x86-64, meaning target-independent compiler options should behave consistently across the two platforms. The intent and expectation is that makefiles and build scripts used to drive PGI compilers on Linux x86-64 should work with little or no modification on OpenPOWER. The `-help` compiler option lists all compiler options by default, or can be used to check the functionality of a specific option. For example:

```
% pgcc -help -fast
Reading rcfile /opt/pgi/linuxpower/17.4/bin/.pgccrc
-fast                Common optimizations; includes -O2 -Munroll=c:1 -Mlre -
Mautoinline          == -Mvect=simd -Mflushz
-M[no]vect[=[no]simd|[no]assoc|[no]fuse]
                    Control automatic vector pipelining
                    Generate [don't generate] SIMD instructions
    [no]simd
    [no]assoc        Allow [disallow] reassociation
    [no]fuse         Enable [disable] loop fusion
-M[no]flushz        Set SSE to flush-to-zero mode
%
```

2.4. Fortran Language

The PGFORTRAN compiler supports Fortran 2003 features as defined in the document ISO/IEC 1539-1 : 2004, *Information technology – Programming Languages – Fortran*, Geneva, 2004 (Fortran 2003).

2.5. C Language

The PGCC compiler supports C99 and many of the important C11 language features as defined in the document ISO/IEC 9899:2013, *Information Technology – Programming Languages – C*, Geneva, 2011 (C11). In particular, thread-local storage, type generics, the `__Noreturn` specifier and `__Alignof`/`__Alignas` are all supported. Certain C11 features

including anonymous **structs/unions**, **_static_assert**, atomics and unicode are not yet supported. PGCC supports compiler options `-c89`, `-c99` and `-c11` to enable/restrict support for/to C89, C99 and C11 respectively. Support for C99 is default. The `-c11` compiler option must be used to enable support and processing of the C11 features.

2.6. C++ Language

The PGC++ compiler supports C++14 as defined in the document ISO/IEC 14882:2014, *Information Technology – Programming Languages – C++*, Geneva. The `--c++14` compiler option must be used to enable support and processing of C++14 features, and the `--c++11` compiler option must be used to enable support and processing of C++11 features. PGC++ is substantially compatible with GNU `g++` through GCC 6.2. In particular it uses GNU name-mangling conventions, uses GCC header files and libraries directly from a standard GCC installation, and is designed to be link-compatible with `g++`.

2.7. OpenACC

The `pgcc`, `pgc++` and `pgfortran` compilers implement OpenACC 2.5 as defined in *The OpenACC Application Programming Interface*, Version 2.5, August 2013, <http://www.openacc.org>, with the exception that these features are not yet supported:

- ▶ nested parallelism
- ▶ declare link
- ▶ enforcement of the new **cache** clause restriction that all references to listed variables must lie within the region being cached

With respect to OpenACC and Tesla support, the PGI 17.4 compilers for OpenPOWER include the same features as the PGI 17.4 compilers for x86-64; these features are described in Chapter 7 of the *PGI User's Guide for OpenPOWER* (<http://www.pgroup.com/doc/pgi16ug-openpower.pdf>). Likewise, as pertains to CUDA Fortran and NVIDIA Tesla GPUs, all features documented in the *PGI CUDA Fortran Programming Guide and Reference* are all functional in the Fortran compiler (<http://www.pgroup.com/doc/pgicudaforug.pdf>), including Beta support for the **managed** attribute and CUDA Unified Memory.

The PGI 17.4 compiler installation for OpenPOWER includes by default the PGI OpenACC Unified Memory Evaluation Package, which is a Beta feature in PGI 15.4 and later production PGI Accelerator OpenACC compilers on x86-64. The package includes header files, object files and a library. The compiler option to enable CUDA Unified Memory is `-ta=tesla:managed`. It is a compile- and link-time option which causes all C/C++ **malloc/calloc/free** dynamic memory allocations, as well as C++ **new/delete**, to be performed using CUDA Unified Memory; similarly, all Fortran **ALLOCATE/DEALLOCATE** statements are performed using CUDA Unified Memory.

When an OpenACC executable is compiled with `-ta=tesla:managed`, the OpenACC runtime dynamically checks variables that normally would be moved to/from accelerator memory. If they are allocated in CUDA Unified Memory the runtime performs no data movement on these variables, allowing the CUDA driver to manage

any necessary data movement between system and device memories. Otherwise, the OpenACC runtime proceeds normally. OpenACC runtime-managed data and dynamically-allocated driver-managed CUDA Unified Memory can be used in the same program. For a more extensive description of the capabilities of this feature, see <http://www.pgroup.com/lit/articles/insider/v6n2a4.htm>.

Chapter 3.

INSTALLATION AND CONFIGURATION

Follow these steps to install PGI 17.4 compilers on an OpenPOWER system. The default installation directory is `/opt/pgi`, but it can be any directory:

```
% tar xzpf pgilinux-2017-174-ppc64le.tar.gz
% ./install
<answer installation questions, assent to licenses>
...
```

Typically for this release, you will want to choose the following during the installation:

1. Choose a "Single-system install", not a "Network install".
2. Install the PGI software in the default `/opt/pgi` directory.
3. Install the CUDA toolkit.

This installs CUDA components in the PGI directory tree, and will not affect a standard CUDA installation on the same system in any way.

4. Install the OpenACC Unified Memory Evaluation package.
5. Create links in the 2017 directory.

This is the directory where CUDA is installed, along with example programs; links are created to the subdirectories of `/opt/pgi/linuxpower/17.4`.

6. Install Open MPI.

3.1. License Management

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

If you purchased a perpetual license and have obtained your new license key, either replace the contents of `license.pgi` with your new license key, or set the environment variable `LM_LICENSE_FILE` to the full path of the desired license file.

If you have not yet obtained your new license key, please consult your PGI order confirmation email for instructions for obtaining and installing your permanent license key. Contact PGI Sales at sales@pgroup.com if you need assistance.

Usage Logging: This release provides per-user records of most recent use in the `.pgiusage` subdirectory inside the main installation directory. Set the environment variable `PGI_LOG_DIRECTORY` to specify a different directory for usage logging.

3.2. Environment Initialization

Assuming the software is installed in `/opt/pgi`, use these commands in `csh` to initialize your environment for use of the PGI compilers:

```
% setenv PGI /opt/pgi
% setenv MANPATH "$MANPATH:$PGI/linuxpower/2017/man"
% set path=($PGI/linuxpower/2017/bin $path)
% which pgc++
/opt/pgi/linuxpower/2017/bin/pgc++
%
```

In `bash`, `sh` or `ksh`, use these commands:

```
% export PGI=/opt/pgi
% export MANPATH=$MANPATH:$PGI/linuxpower/2017/man
% export PATH=$PGI/linuxpower/2017/bin:$PATH
% which pgc++
/opt/pgi/linuxpower/2017/bin/pgc++
%
```

The PGI directory structure is designed to accommodate co-installation of multiple PGI versions. When 17.4 is installed, it will be installed by default in the directory `/opt/pgi/linuxpower/17.4` and links can optionally be created to its sub-directories to make 17.4 default without affecting a previous (e.g., 16.10) install. Non-default versions of PGI compilers that are installed can be used by specifying the `-V<ver>` option on the compiler command line.

3.3. Network Installations

PGI compilers for OpenPOWER may be installed locally on each machine on a network or they may be installed once on a shared file system available to every machine. With the shared file system method, after the initial installation you can run a simple script on each machine to add that system to the family of machines using the common compiler installation. Using this approach, you can create a common installation that works on multiple linuxpower systems even though each system may have different versions of `gcc/libc`.

Follow these steps to create a shared file system installation on OpenPOWER systems:

1. Create a commonly-mounted directory accessible to every system using the same directory path (for example, `/opt/pgi`).
2. Define a locally-mounted directory with a pathname that is identical on all systems. That is, each system has a local directory path with the same pathname (for example `/local/pgi/17.4/share_objects`). Runtime libraries which are *libc*-version dependent will be stored here. This will ensure that executable files built on one system will work on other systems on the same network.
3. Run the install script for the first installation:

```
% tar xzpf pgilinux-2017-174-ppc64le.tar.gz
% ./install
<answer installation questions, assent to licenses>
...
```

At the "Please choose install option:" prompt, choose "Network install".

4. Once the initial PGI installation is complete, configure the environment as described in the preceding section.
5. On each subsequent system, follow these steps:
 - a. Set up the environment as described in the preceding section.
 - b. Run the `add_network_host` script which is now in your `$PATH`:

```
$ add_network_host
```

and the compilers should now work.

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, visit the frequently asked questions (FAQ) section on pgroup.com at <http://www.pgroup.com/support/faq.htm>.

4.1. Release Specific Limitations

The following PGI features are limited or are not implemented in the 17.4 release for OpenPOWER+Tesla:

- ▶ -Mipa is not enabled (no PGI inter-procedural analysis/optimization); the command-line option is accepted and silently ignored.
- ▶ -Mpf/-Mpfo are not enabled (no profile-feedback optimization); the command-line options are accepted and silently ignored.
- ▶ No debugging support for Fortran.

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: sales@pgroup.com

WWW: <http://www.pgroup.com>

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>

Many questions and problems can be resolved by following instructions and the information available at our frequently asked questions (FAQ) site:

<http://www.pgroup.com/support/faq.htm>

Submit technical support requests through the online form at:

https://www.pgroup.com/support/support_request.php

PGI documentation is available at <http://www.pgroup.com/resources/docs.htm>.

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

PGI Workstation, PGI Server, PGI Accelerator, PGF95, PGF90, PGFORTRAN, and PGI Unified Binary are trademarks; and PGI, PGHPF, PGF77, PGCC, PGC++, PGI Visual Fortran, PVF, PGI CDK, Cluster Development Kit, PGPROF, PGDBG, and The Portland Group are 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

© 2013-2017 NVIDIA Corporation. All rights reserved.