

PVF Release Notes

Version 2017



PGI Compilers and Tools

TABLE OF CONTENTS

Chapter 1. PVF Release Overview.....	1
1.1. Product Overview.....	1
1.2. Microsoft Build Tools.....	2
1.3. Terms and Definitions.....	2
Chapter 2. New and Modified Features.....	3
2.1. What's New in Release 2017.....	3
2.2. New and Modified Compiler Options.....	5
2.3. New and Modified Fortran Functionality.....	5
2.4. Updates to CUDA Toolkit Support.....	5
2.5. New Features in PGI Accelerator OpenACC Compilers.....	6
2.6. Runtime Library Routines.....	7
Chapter 3. Selecting an Alternate Compiler.....	8
3.1. For a Single Project.....	8
3.2. For All Projects.....	8
Chapter 4. Distribution and Deployment.....	10
4.1. Application Deployment and Redistributables.....	10
4.1.1. PGI Redistributables.....	10
4.1.2. Microsoft Redistributables.....	10
Chapter 5. Troubleshooting Tips and Known Limitations.....	12
5.1. PVF IDE Limitations.....	12
5.2. PVF Debugging Limitations.....	12
5.3. PGI Compiler Limitations.....	13
5.4. OpenACC Issues.....	13
5.5. Corrections.....	13
Chapter 6. Contact Information.....	14

Chapter 1.

PVF RELEASE OVERVIEW

Welcome to Release 2017 of PGI Visual Fortran®, a set of Fortran compilers and development tools for Windows integrated with Microsoft® Visual Studio.

This document describes the new features of the PVF IDE interface, differences in the PVF 2017 compilers and tools from previous releases, and late-breaking information not included in the standard product documentation.



PGI Release 2016 version 16.4 and newer includes FlexNet license daemons updated to version 11.13.1.3. This update addresses a [FlexNet security vulnerability](#). These new license daemons also work with older PGI releases. We recommend all users update their license daemons - see the [licensing FAQ](#) for more information. This FlexNet update also requires you to [update your PGI FlexNet license keys](#) to a new format. Older keys are incompatible.

1.1. Product Overview

PVF is integrated with two versions of Microsoft Visual Studio. Currently, Visual Studio 2013 and 2015 are supported. Throughout this document, "PGI Visual Fortran" refers to PVF integrated with either of the two supported versions of Visual Studio. Similarly, "Microsoft Visual Studio" refers to Visual Studio 2013 and VS 2015. When it is necessary to distinguish among the products, the document does so.

Single-user node-locked and multi-user network floating license options are available for both products. When a node-locked license is used, one user at a time can use PVF on the single system where it is installed. When a network floating license is used, a system is selected as the server and it controls the licensing, and users from any of the client machines connected to the license server can use PVF. Thus multiple users can simultaneously use PVF, up to the maximum number of users allowed by the license.

PVF provides a complete Fortran development environment fully integrated with Microsoft Visual Studio. It includes a custom Fortran Build Engine that automatically derives build dependencies, Fortran extensions to the Visual Studio editor, a custom PGI Debug Engine integrated with the Visual Studio debugger, PGI Fortran compilers, and PVF-specific property pages to control the configuration of all of these.

Release 2017 of PGI Visual Fortran includes the following components:

- ▶ PGFORTRAN OpenMP and auto-parallelizing Fortran 2003 compiler.
- ▶ PGF77 OpenMP and auto-parallelizing FORTRAN 77 compiler.
- ▶ PVF Visual Studio integration components.
- ▶ OpenACC and CUDA Fortran tools and libraries necessary to build executables for Accelerator GPUs, when the user's license supports these optional features.
- ▶ PVF documentation.

If you do not already have Microsoft Visual Studio on your system, be sure to get the PVF installation package that contains the Visual Studio 2015 shell.

1.2. Microsoft Build Tools

PVF on all Windows systems includes Microsoft Open Tools. These files are required in addition to the files Microsoft provides in the Windows SDK.

1.3. 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 online glossary at <http://www.pgroup.com/support/definitions.htm>.

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 and Haswell 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, and AMD Piledriver processors.

Chapter 2.

NEW AND MODIFIED FEATURES

This section provides information about the new and modified features of Release 2017 of PGI Visual Fortran.

2.1. What's New in Release 2017

17.1 Updates and Additions



IMPORTANT: The PGI Visual Fortran 2017 release no longer supports 32-bit development. PVF 2017 only supports 64-bit operating systems and does not include the ability to compile 32-bit applications for execution on either 32-bit or 64-bit operating systems. With PGI 2017, you can still view the Win32 solution platform within a PVF project, but you cannot build or debug using the Win32 solution platform. If you have not yet migrated your projects from the Win32 solution platform to the x64 solution platform, consider doing so by using PVF 16.10 or earlier versions which support building both the Win32 and x64 versions for comparison.

- ▶ All PGI Compilers
 - ▶ Updated how floating point divides are computed to guarantee results will be uniform for scalar and vector operations. This change can cause numerical differences between PGI 17.1 and previous PGI releases. In rare cases, this change can cause an increase in execution time.
 - ▶ 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 control inlining, 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.
- ▶ PGI Accelerator OpenACC Compilers
 - ▶ 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.
 - ▶ Removed compute capability 2.0 (Fermi) from the compilers' default set of compute capabilities. The `cc20` sub-option to the `-ta=tesla` and `-Mcuda` options is deprecated.
 - ▶ Improved data management within 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 references to variables within a cached region, OpenACC 2.5 feature support is complete in PGI 2017.

- ▶ Other Features and Additions
 - ▶ Added support for Microsoft Windows Server 2016.
- ▶ Deprecations and Eliminations
 - ▶ PGI 2017 supports 64-bit operating systems only. Compiling 32-bit applications for execution on either 32-bit or 64-bit operating systems is no longer supported on any platform.
 - ▶ Dropped support for the CUDA 7.0 toolkit.

2.2. New and Modified Compiler Options

Release 2017 supports new and updated command line options and keyword suboptions.

Added the following options:

- ▶ `-cpp` is now an alias for `-Mpreprocess`.
- ▶ `[dis]allows` variadic macros.

Changed the following `-Minline` sub-options:

- ▶ Added `totalsize:n` to limit inlining to the total size of `n`.
- ▶ `maxsize:n` replaces `size:n` to prevent inlining of functions bigger than `n`. The compilers silently convert the previous `size:n` to `maxsize:n`.
- ▶ Removed `levels:n` which limited inlining to `n` levels of functions. The compilers silently ignore `levels:n`.

2.3. New and Modified Fortran Functionality

- ▶ TBD.

2.4. Updates to CUDA Toolkit Support

As of PGI 17.1, CUDA 7.5 and the CUDA 8.0 are supported on Linux and Windows. CUDA 7.5 is the default.

Support for PGI Accelerator features for macOS including CUDA Fortran, OpenACC and CUDA-x86 is no longer available as of the PGI 2017 release.

Targeting a CUDA Toolkit Version

- ▶ The CUDA 7.5 Toolkit is set as the default in PGI 17.1. 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 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. To use the CUDA 8.0 Toolkit, you must download and install the CUDA 8.0 driver from NVIDIA.
- ▶ `pgaccelinfo` prints the driver version as the first line of output. For a 7.5 driver, it prints:

```
CUDA Driver Version 7050
```


2.5. New Features in PGI Accelerator OpenACC Compilers

Multicore Support

PGI Accelerator OpenACC compilers support the option `-ta=multicore`, to set the target accelerator for OpenACC programs to the host multicore. This will compile OpenACC compute regions for parallel execution across the cores of the host X86 processor or processors. The host multicore will be treated as a shared-memory accelerator, so the data clauses (`copy`, `copyin`, `copyout`, `create`) will be ignored and no data copies will be executed.

By default, `-ta=multicore` will generate code that will use all the available cores of the processor. If the compute region specifies a value in the `num_gangs` clause, the minimum of the `num_gangs` value and the number of available cores will be used. At runtime, the number of cores can be limited by setting the environment variable `ACC_NUM_CORES` to a constant integer value. If an OpenACC compute construct appears lexically within an OpenMP parallel construct, the OpenACC compute region will generate sequential code. If an OpenACC compute region appears dynamically within an OpenMP region or another OpenACC compute region, the program may generate many more threads than there are cores, and may produce poor performance.

The `ACC_BIND` environment variable is set by default with `-ta=multicore`; `ACC_BIND` has similar behavior to `MP_BIND` for OpenMP.

The `-ta=multicore` option differs from the `-ta=host` option in that `-ta=host` generates sequential code for the OpenACC compute regions. In this release, `-ta=multicore` is not supported in conjunction with `-ta=tesla` or `-ta=host`.

Default Compute Capability

The default compute capability list for OpenACC and CUDA Fortran compilation for NVIDIA Tesla targets is `cc30, cc35, and cc50`. If CUDA 8.0 is specified using the `cuda8.0` sub-option, `cc60` is added to the default compute capability list. The generation of device code can be time consuming, so you may notice an increase in compile time. You can override the default by specifying one or more compute capabilities using either command-line options or an `rcfile`.

To change the default with a command-line option, provide a comma-separated list of compute capabilities to `-ta=tesla:` for OpenACC or `-Mcuda=` for CUDA Fortran.

To change the default with an `rcfile`, set the `DEFCOMPUTECAP` value to a blank-separated list of compute capabilities in the `siterc` file located in your installation's `bin` directory:

```
set DEFCOMPUTECAP=30 35 50;
```

Alternatively, if you don't have permissions to change the `siterc` file, you can add the `DEFCOMPUTECAP` definition to a separate `.mypgirc` file (`mypgi_rc` on Windows) in your home directory.

OpenACC 2.5 Features Not Supported

The PGI 2017 release contains support for all OpenACC 2.5 features except adding restrictions to `cache` clause variable refs to variables within a cached region.

OpenACC 2.0 Features Not Supported

The PGI 2017 release contains support for all OpenACC 2.0 features except the following:

- ▶ The `declare link` directive for global data is not implemented.
- ▶ Nested parallelism (`parallel` and `kernels` constructs within a `parallel` or `kernels` region) is not implemented.

Support for Profiler/Trace Tool Interface

PGI compilers support the OpenACC 2.5 version of the OpenACC Profiler/Trace Tools Interface. This is the interface used by the PGI profiler to collect performance measurements of OpenACC programs.

2.6. Runtime Library Routines

PGI 2017 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.

Chapter 3.

SELECTING AN ALTERNATE COMPILER

Each release of PGI Visual Fortran contains two components — the newest release of PVF and the newest release of the PGI compilers and tools that PVF targets.

When PVF is installed onto a system that contains a previous version of PVF, the previous version of PVF is replaced. The previous version of the PGI compilers and tools, however, remains installed side-by-side with the new version of the PGI compilers and tools. By default, the new version of PVF will use the new version of the compilers and tools. Previous versions of the compilers and tools may be uninstalled using Control Panel | Add or Remove Programs.

There are two ways to use previous versions of the compilers:

- ▶ Use a different compiler release for a single project.
- ▶ Use a different compiler release for all projects.

The method to use depends on the situation.

3.1. For a Single Project

To use a different compiler release for a single project, you use the compiler flag `-V<ver>` to target the compiler with version `<ver>`. This method is the recommended way to target a different compiler release.

For example, `-V13.8` causes the compiler driver to invoke the 13.8 version of the PGI compilers if these are installed.

To use this option within a PVF project, add it to the Additional options section of the Fortran | Command Line and Linker | Command Line property pages.

3.2. For All Projects

You can use a different compiler release for all projects.

The Tools | Options dialog within PVF contains entries that can be changed to use a previous version of the PGI compilers. Under Projects and Solutions | PVF

Directories, there are entries for Executable Directories, Include and Module Directories, and Library Directories.

- ▶ For the x64 platform, each of these entries includes a line containing `$(PGIToolsDir)`. To change the compilers used for the x64 platform, change each of the lines containing `$(PGIToolsDir)` to contain the path to the desired `bin`, `include`, and `lib` directories.



Warning: The debug engine in PVF 2017 is not compatible with previous releases. If you use `Tools | Options` to target a release prior to 2017, you cannot use PVF to debug. Instead, use the `-v` method described earlier in this section to select an alternate compiler.

Chapter 4.

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.

4.1. 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.

4.1.1. PGI Redistributables

PGI Visual Fortran includes redistributable directories which contain all of the PGI dynamically linked libraries that can be re-distributed by PVF 2017 licensees under the terms of the PGI End-User License Agreement (EULA). For reference, a copy of the PGI EULA in PDF form is included in the release.

The following paths for the redistributable directories assume 'C:' is the system drive.

- ▶ On a 64-bit Windows system, the redistributable directory is:

```
C:\Program Files\PGI\win64\17.1\REDIST
```

The redistributable directories contain the PGI runtime library DLLs for all supported targets. This enables users of the PGI compilers to create packages of executables and PGI runtime libraries that execute successfully on almost any PGI-supported target system, subject to the requirement that end-users of the executable have properly initialized their environment to use the relevant version of the PGI DLLs.

4.1.2. Microsoft Redistributables

PGI Visual Fortran includes Microsoft Open Tools, the essential tools and libraries required to compile, link, and execute programs on Windows. PVF 2017 installed for Microsoft Visual Studio 2015 includes the latest version, version 14.0, of the Microsoft Open Tools; PVF 2017

installed for Microsoft Visual Studio 2013 includes the previous version, version 12.0, of the Microsoft Open Tools in order to maintain compatibility with Visual C++ 2013.

The Microsoft Open Tools directory contains a subdirectory named `REDIST`. PGI 2017 licensees may redistribute the files contained in this directory in accordance with the terms of the associated license agreements.



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

Chapter 5.

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>.

5.1. PVF IDE Limitations

The issues in this section are related to IDE limitations.

- ▶ When moving a project from one drive to another, all `.d` files for the project should be deleted and the whole project should be rebuilt. When moving a solution from one system to another, also delete the solution's Visual Studio Solution User Options file (`.suo`).
- ▶ The Resources property pages are limited. Use the Resources | Command Line property page to pass arguments to the resource compiler. Resource compiler output must be placed in the intermediate directory for build dependency checking to work properly on resource files.
- ▶ Dragging and dropping files in the Solution Explorer that are currently open in the Editor may result in a file becoming "orphaned." Close files before attempting to drag-and-drop them.

5.2. PVF Debugging Limitations

The following limitations apply to PVF debugging:

- ▶ Debugging of unified binaries is not fully supported. The names of some subprograms are modified in the creation of the unified binary, and the PVF debug engine does not translate these names back to the names used in the application source code. For more information on debugging a unified binary, refer to www.pgroup.com/support/tools.htm.
- ▶ In some situations, using the Watch window may be unreliable for 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.

- ▶ Rolling over Fortran arrays during a debug session is not supported when Visual Studio is in Hex mode. This limitation also affects Watch and Quick Watch windows.

Workaround: deselect Hex mode when rolling over arrays.

5.3. PGI Compiler Limitations

- ▶ Take extra care when using `-Mprof` with PVF runtime library DLLs. To build an executable for profiling, use of the static libraries is recommended. The static libraries are used by default in the absence of `-Bdynamic`.
- ▶ Using `-Mpf` and `-mp` together is not supported. The `-Mpf` flag disables `-mp` at compile time, which can cause runtime 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. Using the `-Mpf0` flag does not disable OpenMP processing.

5.4. 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.

5.5. Corrections

A number of problems are corrected in this release. Refer to www.pgroup.com/support/release_tprs.htm for a complete and up-to-date table of technical problem reports, TPRs, fixed in recent releases of the PGI compilers and tools. This table contains a summary description of each problem as well as the version in which it was fixed.

Chapter 6.

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.

PGI[®]