

PGI[®] COMPILERS & TOOLS

PGI VISUAL FORTRAN RELEASE NOTES

Version 2018



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. What's New in PGI 2018.....	3
2.1. What's New in 18.5.....	3
2.2. What's New in 18.4.....	3
2.3. What's New in 18.3.....	4
2.4. What's New in 18.1.....	4
2.5. New and Modified Compiler Options.....	7
2.6. CUDA Toolkit Versions.....	8
2.7. OpenMP.....	9
2.8. Runtime Library Routines.....	10
Chapter 3. Selecting an Alternate Compiler.....	11
3.1. For a Single Project.....	11
3.2. For All Projects.....	11
Chapter 4. Distribution and Deployment.....	13
4.1. Application Deployment and Redistributables.....	13
4.1.1. PGI Redistributables.....	13
4.1.2. Microsoft Redistributables.....	13
Chapter 5. Troubleshooting Tips and Known Limitations.....	15
5.1. PVF IDE Limitations.....	15
5.2. PVF Debugging Limitations.....	15
5.3. PGI Compiler Limitations.....	16
5.4. OpenACC Issues.....	16
Chapter 6. Contact Information.....	17

LIST OF TABLES

Table 1	Third-Party Software Security Updates for PGI version 18.5	7
---------	--	---

Chapter 1.

PVF RELEASE OVERVIEW



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 in What's New in PGI 2018](#) and our [FlexNet Update FAQ](#) for more information.

This chapter provides an overview of Release 2018 of PGI Visual Fortran[®], a set of Fortran compilers and development tools for Windows integrated with Microsoft[®] Visual Studio.

1.1. Product Overview

PVF is integrated with Microsoft Visual Studio 2015. Throughout this document, "PGI Visual Fortran" refers to PVF integrated with VS 2015. Similarly, "Microsoft Visual Studio" refers to Visual Studio 2015. When it is necessary to distinguish further, 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 2018 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.

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 [PGI online glossary](http://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.

Chapter 2.

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 PGI Visual Fortran!

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.

2.1. What's New in 18.5

The PGI 18.5 release contains all features and fixes found in PGI 18.4 and a few key updates for important user-reported problems.

Added full support for CUDA 9.2; use the `cuda9.2` sub-option with the `-ta=tesla` or `-Mcuda` compiler options to compile and link with the integrated CUDA 9.2 toolkit components.

Added support for Xcode 9.3 on macOS.

Improved function offset information in runtime tracebacks in optimized (non-debug) modes.

2.2. What's New in 18.4

The PGI 18.4 release contains all features and fixes found in PGI 18.3 and a few key updates for important user-reported problems.

Added support for internal procedures, assigned procedure pointers and internal procedures passed as actual arguments to procedure dummy arguments.

Added support for intrinsics SCALE, MAXVAL, MINVAL, MAXLOC and MINLOC in initializers.

2.3. What's New in 18.3

The PGI 18.3 release contains all the new features found in PGI 18.1 and a few key updates for important user-reported problems.

2.4. 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.

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.

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`, `libpgftrntl`, and `libpgf90rtl` to `libpgm`. This change should be transparent unless you have been explicitly adding `libpgc`, `libpgftrntl`, 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.

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 Microsoft Visual Studio 2013. Installing PVF 18.5 will cause any version of PVF integrated with VS 2013 to stop working. PVF continues to support VS 2015.

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 `-Mcudax86` 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 18.5

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 .

2.5. New and Modified Compiler Options

Release 2018 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`.

2.6. 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 18.5 contains the following versions of the CUDA Toolkits:

- ▶ CUDA 8.0 (default)
- ▶ CUDA 9.0
- ▶ CUDA 9.1
- ▶ CUDA 9.2

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 9.2 Toolkit you would use:

```
pgcc -ta=tesla:cuda9.2
```

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 9.2 Toolkit as the default, add the following line to one of these files:

```
set DEFCUDAVERSION=9.2;
```

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.

The ability to compile with a CUDA Toolkit other than the versions installed with the PGI compilers is a feature not currently supported on Windows.

To use a CUDA toolkit that is not installed with a PGI release, such as CUDA 7.5 with PGI 18.5, 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. OpenMP

OpenMP 3.1

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

2.8. 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.

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 2018 is not compatible with previous releases. If you use `Tools | Options` to target a release prior to 2018, 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 2018 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\18.5\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 2018 installed for Microsoft Visual Studio 2015 includes version 14.0 of the 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 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, please see the [PGI frequently asked questions \(FAQ\)](#) webpage.

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.

- ▶ 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 `-Mpfi` and `-mp` together is not supported. The `-Mpfi` 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 `-Mpfo` 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.

Chapter 6. CONTACT INFORMATION

You can contact PGI at:

9030 NE Walker Road, Suite 100
Hillsboro, OR 97006

Or electronically using any of the following means:

Fax: +1-503-682-2637

Sales: sales@pgroup.com

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

The [PGI User Forum](http://pgicompile.com/userforum), pgicompile.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](#), pgicompile.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 \(FAQ\)](#), pgicompile.com/faq.

Submit support requests using the [PGI Technical Support Request form](#), pgicompile.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.