Statement of Software Standards Support

05 January 2017

Mr. Reinhard Budich
On behalf of the ENES HPC Task Force
Max Planck Institute for Meteorology
Strategic IT Partnerships
Room 321
Bundesstr. 53
D-20146 Hamburg
Germany

Dear Mr. Budich,

Thank you for the opportunity provided by the ENES HPC Task Force for NVIDIA to proclaim the company's support plans for software standards in the development of HPC and GPU technology. NVIDIA would like to express our general interest in developing high quality HPC solutions for the Earth system modeling community, and specifically this letter describes the status and directions of developments for a comprehensive HPC programming environment.

NVIDIA Corporation headquartered in Santa Clara, CA, is a leader in HPC and computer graphics, with more than 10,000 employees in 20 countries world-wide. With the introduction of the general purpose Graphics Processing Unit (GPU) in 2004, NVIDIA set a milestone in parallel computing. NVIDIA expertise developed in programmable GPUs has led to several breakthroughs in efficient and cost-effective HPC solutions. HPC is an established strategic initiative within the company, to help industry and research address the most complex scientific and engineering challenges.

For the question posed to NVIDIA by the ENES HPC Task Force:

> "In which directions does NVIDIA plan to develop FORTRAN and other standards, and how it will interact with the/our community, and with the standardisation activities/bodies?"

NVIDIA presents the following 4 proclamations:

1. Standards Commitment: NVIDIA and PGI (acquired during Jul 2013) support Fortran/C/C++ on x86 CPU-based systems, and with OpenACC extensions to enable the parallelization and offload of data and computations to NVIDIA GPUs, and/or parallelization across multiple x86 CPU cores. In addition, the PGI Fortran compiler supports CUDA Fortran extensions, which are a direct analog to CUDA C/C++. During Nov 2016, support for these same products and features became available for OpenPOWER CPU-based systems with the 16.10 release of PGI compilers. Additionally for OpenPOWER, CUDA Fortran support is now available from IBM in their XLF compiler since their V15.1.4 release:
http://www.ibm.com/support/knowledgecenter/SSAT4T_15.1.4/com.ibm.xlf1514.lelinux.doc/getstart/cudaf_v1514.html

NVIDIA expects to support all such compiler products long-term, with development of incremental enhancements based on the evolution of Fortran/C/C++ standards, OpenACC and OpenMP specifications, and the CUDA roadmap.

2. Programming Models: NVIDIA programming model directions are well aligned with the requirements of ESM developers: directives-based compiler standards that permit performance and maintenance portability across different hardware systems. The company's leading HPC programming model for the ESM community is OpenACC, and NVIDIA continues with heavy investments in PGI to expand the capabilities and platform support for OpenACC.

During Nov 2016, the PGI Community Edition program was launched: a no-cost license to the current release of the PGI Fortran, C and C++ compilers and tools for multicore CPUs and NVIDIA GPUs, including all OpenACC, OpenMP and CUDA Fortran features. Details are provided in this FAQ from PGI: https://www.pgroup.com/support/community-faq.htm

NVIDIA also contributes to the OpenMP Architectural Review Board (ARB) and other developments of industry standards. For OpenMP, NVIDIA is an active member in all aspects with technical representation led by Jeff Larkin jlarkin@nvidia.com, an engineer in NVIDIA's US Developer Technology Group, and NVIDIA will continue to work within the OpenMP committee to ensure the standard becomes a viable option for GPUs. In addition and since NVIDIA's Statement provided to the ENES during Apr 2016, NVIDIA's James Beyer jbeyer@nvidia.com became co-chair of the OpenMP subcommittee on accelerators.

3. FORTRAN Committee: NVIDIA is an active member of J3 www.j3-fortran.org, the US-based Fortran standards committee, known formally as PL22.3, a technical subcommittee of the International Committee for Information Technology Standards (INCITS) formerly known as the National Committee for Information Technology Standards (NCITS). NVIDIA and PGI have two technical professionals participating since Jan 2015, Mark LeAir and Bob Scollard, at least one of whom attends all meetings of the J3.

Within J3, NVIDIA holds 1 of 9 positions as VOTING PRINCIPAL: http://www.j3-fortran.org/doc/standing/links/001.txt
- Mark LeAir, NVIDIA Corp, Beaverton, OR, USA; mark.leair@pgroup.com; +1 (503) 748-6463; Alternate: Bob Scollard

Generally NVIDIA and PGI have a goal to track more closely to the current version of the Fortran standard than PGI was able to achieve prior to the PGI acquisition by NVIDIA. NVIDIA's participation is motivated both by this goal, and by the company vision to enable programming of heterogeneous CPU+GPU systems using standard languages.

4. FORTRAN Contributions: Further confirmation of NVIDIA and PGI commitments to Fortran advancement is a US Department of Energy (DOE) program announced during Nov 2015 under which PGI is funded to develop an open-source Fortran compiler designed for integration with the widely used LLVM compiler infrastructure. Source code for the resulting Fortran front-end is expected to be available in open-source form during late Q1 2017.

The DOE announcement is available at:
https://www.llnl.gov/news/nnsa-national-labs-team-nvidia-develop-open-source-fortran-compiler-technology

A FAQ of the project is available at: https://groups.google.com/forum/#!topic/llvm-dev/fLfOQKNEWBU

Indeed NVIDIA would greatly benefit from the deep knowledge and experiences of the ENES, and we invite valuable feedback from the ENES HPC Task Force, and the potential for technical collaboration with the ENES community. Thank you again for this opportunity, and please contact me with questions ENES may have regarding this statement.

Sincerely,

Stan Posey
HPC Program Manager
ESM and CFD Solutions Development
NVIDIA Corporation
408-207-7439
sposey@nvidia.com

With contributions from others at NVIDIA/PGI:
- Doug Miles, Director, PGI Compilers and Tools
- Duncan Poole, NVIDIA, President of the OpenACC Executive Committee
- Jeff Larkin, Senior Engineer, Developer Technology Group, NVIDIA Representative to the OpenMP ARB