



NVIDIA HPC SDK RELEASE NOTES

RN-09976-001-V22.2 | February 2022



TABLE OF CONTENTS

- Chapter 1. What's New..... 1
- Chapter 2. Release Component Versions..... 2
- Chapter 3. Supported Platforms..... 4
 - 3.1. Platform Requirements for the HPC SDK..... 4
 - 3.2. Supported CUDA Toolchain Versions..... 5
- Chapter 4. Known Limitations..... 6
- Chapter 5. Deprecations and Changes..... 7

LIST OF TABLES

Table 1 HPC SDK Release Components	2
Table 2 HPC SDK Platform Requirements	4

Chapter 1.

WHAT'S NEW

Welcome to the 22.2 version of the NVIDIA HPC SDK, a comprehensive suite of compilers and libraries enabling developers to program the entire HPC platform, from the GPU foundation to the CPU and out through the interconnect.

Some key changes that are new in this release of the NVIDIA HPC SDK for Linux include:

- ▶ Support for array reductions has been significantly improved for directives-based programming models.
- ▶ 128-bit integers are available by default.

Chapter 2.

RELEASE COMPONENT VERSIONS

The NVIDIA HPC SDK 22.2 release contains the following versions of each component:

Table 1 HPC SDK Release Components

	Linux_x86_64				Linux_ppc64le				Linux_aarch64			
	CUDA 10.2	CUDA 11.0	CUDA 11.2	CUDA 11.6	CUDA 10.2	CUDA 11.0	CUDA 11.2	CUDA 11.6	CUDA 10.2	CUDA 11.0	CUDA 11.2	CUDA 11.6
nvc++	22.2				22.2				22.2			
nvc	22.2				22.2				22.2			
nvfortr	22.2				22.2				22.2			
nvcc	10.2.89	11.0.228	11.2.146	11.6.55	10.2.89	11.0.228	11.2.146	11.6.55	N/A	11.0.228	11.2.146	11.6.55
NCCL	2.11.4	2.11.4	2.11.4	2.11.4	2.11.4	2.11.4	2.11.4	2.11.4	N/A	2.11.4	2.11.4	2.11.4
NVSHM	2.4.1	2.4.1	2.4.1	2.4.1	2.4.1	2.4.1	2.4.1	2.4.1	N/A	N/A	N/A	N/A
cuBLAS	0.2.2.8	11.2.0.25	11.4.1.10	11.8.1.74	0.2.2.8	11.2.0.25	11.4.1.10	11.8.1.74	N/A	11.2.0.25	11.4.1.10	11.8.1.74
cuFFT	10.1.2.8	10.2.1.24	10.4.0.13	10.7.0.5	10.1.2.8	10.2.1.24	10.4.0.13	10.7.0.5	N/A	10.2.1.24	10.4.0.13	10.7.0.5
cuRAND	10.1.2.8	10.2.1.24	10.2.3.13	10.2.9.5	10.1.2.8	10.2.1.24	10.2.3.13	10.2.9.5	N/A	10.2.1.24	10.2.3.13	10.2.9.5
cuSOLV	10.3.0.8	10.6.0.24	11.1.0.13	11.2.1.4	10.3.0.8	10.6.0.24	11.1.0.13	11.2.1.4	N/A	10.6.0.24	11.1.0.13	11.2.1.4
cuSOLV	N/A	N/A	N/A	0.1.0	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
cuSPAR	10.3.1.8	11.1.1.24	11.4.0.13	11.7.1.5	10.3.1.8	11.1.1.24	11.4.0.13	11.7.1.5	N/A	11.1.1.24	11.4.0.13	11.7.1.5
cuTENS	1.4.0	1.4.0	1.4.0	1.4.0	1.4.0	1.4.0	1.4.0	1.4.0	N/A	1.4.0	1.4.0	1.4.0
Nsight Compu	2022.1.0				2022.1.0				2022.1.0			
Nsight System	2021.5.1.118				2021.5.1.118				2021.5.1.118			
OpenM	3.1.5				3.1.5				3.1.5			
HPC-X	N/A	2.10	2.10	2.10	N/A	N/A	N/A	N/A	N/A	2.10	2.10	2.10
UCX	N/A	1.12.0	1.12.0	1.12.0	N/A	N/A	N/A	N/A	N/A	1.12.0	1.12.0	1.12.0

	Linux_x86_64				Linux_ppc64le				Linux_aarch64			
	CUDA 10.2	CUDA 11.0	CUDA 11.2	CUDA 11.6	CUDA 10.2	CUDA 11.0	CUDA 11.2	CUDA 11.6	CUDA 10.2	CUDA 11.0	CUDA 11.2	CUDA 11.6
OpenBLAS	0.3.13				0.3.13				0.3.13			
Scalapy	2.1.0				2.1.0				2.1.0			
Thrust	1.9.7	1.9.9	1.10.0	1.15.0	1.9.7	1.9.9	1.10.0	1.15.0	1.9.7	1.9.10	1.10.0	1.15.0
CUB	N/A	1.9.9	1.10.0	1.15.0	N/A	1.9.9	1.10.0	1.15.0	N/A	1.9.9	1.10.0	1.15.0
libcudpp	N/A	1.0.0	1.3.0	1.7.0	N/A	1.0.0	1.3.0	1.7.0	N/A	1.0.0	1.3.0	1.7.0

Chapter 3. SUPPORTED PLATFORMS

3.1. Platform Requirements for the HPC SDK

Table 2 HPC SDK Platform Requirements

Architecture	Linux Distributions	Minimum gcc/ glibc Toolchain	Minimum CUDA Driver
x86_64	CentOS 7.2, 7.3, 7.4, 7.5, 7.6, 7.7, 7.8 CentOS 7.9, 8.0, 8.1, 8.2 Fedora 29, 30, 31, 32 OpenSUSE Leap 15.0, 15.1 RHEL 7.0, 7.1, 7.2, 7.3, 7.4, 7.5, 7.6, 7.7, 7.8, 7.9 RHEL 8.0, 8.1, 8.2 SLES 12SP4, 12SP5, 15SP1 Ubuntu 18.04, 20.04	C99: 4.8 C11: 4.9 C++03: 4.8 C++11: 4.9 C++14: 5.1 C++17: 7.1	440.33
ppc64le	RHEL 7.3, 7.4, 7.5, 7.6, 7.7, 8.0, 8.1 RHEL Pegas 7.5, 7.6 Ubuntu 18.04	C99: 4.8 C11: 4.9 C++03: 4.8 C++11: 4.9 C++14: 5.1 C++17: 7.1	440.33
aarch64	CentOS 8.1, 8.2, 8.3, 8.4	C99: 4.8 C11: 4.9	450.36

Architecture	Linux Distributions	Minimum gcc/ glibc Toolchain	Minimum CUDA Driver
	RHEL 8.1, 8.2 Ubuntu 18.04, 20.04	C++03: 4.8 C++11: 4.9 C++14: 5.1 C++17: 7.1	

Programs generated by the HPC Compilers for x86_64 processors require a minimum of AVX instructions, which includes Sandy Bridge and newer CPUs from Intel, as well as Bulldozer and newer CPUs from AMD. POWER 8 and POWER 9 CPUs from the POWER architecture are supported. For the Arm architecture, the minimum required version is Arm v8.1.

3.2. Supported CUDA Toolchain Versions

The NVIDIA HPC SDK uses elements of the CUDA toolchain when building programs for execution with NVIDIA GPUs. Every HPC SDK installation package puts the required CUDA components into an installation directory called `[install-prefix]/[arch]/[nvhpc-version]/cuda`.

An NVIDIA CUDA GPU device driver must be installed on a system with a GPU before you can run a program compiled for the GPU on that system. The NVIDIA HPC SDK does not contain CUDA Drivers. You must download and install the appropriate [CUDA Driver from NVIDIA](#), including the [CUDA Compatibility Platform](#) if that is required.

The `nvaccelinfo` tool prints the CUDA Driver version in its output. You can use it to find out which version of the CUDA Driver is installed on your system.

The NVIDIA HPC SDK 22.2 includes the following CUDA toolchain versions:

- ▶ CUDA 10.2
- ▶ CUDA 11.0
- ▶ CUDA 11.2
- ▶ CUDA 11.6

The minimum required CUDA driver versions are listed in the table in Section 3.1.

Chapter 4.

KNOWN LIMITATIONS

- ▶ Some users may experience a bug when using OpenBLAS that causes segmentation faults at job startup; increasing the user's data segment limit (`ulimit -d`) will work around this issue. This issue has recently been addressed upstream, and the fix will be included when OpenBLAS is updated in a future release of the HPC SDK.
- ▶ Debug information for Fortran arrays with non-constant bounds is not handled correctly, and querying values will yield incorrect results. Stepping through CUDA Fortran and OpenACC kernels is partially supported, but incorrect line numbers are displayed. For additional general limitations with `cuda-gdb`, please refer to its documentation.
- ▶ When using `-stdpar` to accelerate C++ parallel algorithms, the algorithm calls cannot include virtual function calls or function calls through a function pointer, cannot use C++ exceptions, can only dereference pointers that point to the heap, and must use random access iterators (raw pointers as iterators work best).
- ▶ When `nvc++ -stdpar=multicore` is used to generate parallel code, OpenMP pragmas in the same translation unit will also be enabled.

Chapter 5. DEPRECATIONS AND CHANGES

- ▶ Starting with the 21.11 version of the NVIDIA HPC SDK, the HPC-X package is no longer shipped as part of the packages made available for the POWER architecture.
- ▶ The current default of `-gpu=implicitsections` will change in a future release to `-gpu=noimplicitsections` to adhere to the OpenACC specification.
- ▶ Starting with the 21.5 version of the NVIDIA HPC SDK, the `-cuda` option for `NVC++` and `NVFORTRAN` no longer automatically links the NVIDIA GPU math libraries. Please refer to the `-cudalib` option.
- ▶ HPC Compiler support for the Kepler architecture of NVIDIA GPUs was deprecated starting with the 21.3 version of the NVIDIA HPC SDK.
- ▶ Support for the KNL architecture of multicore CPUs in the NVIDIA HPC SDK was deprecated in the HPC SDK version 21.3.

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, CUDA, CUDA-X, GPUDirect, HPC SDK, NGC, NVIDIA Volta, NVIDIA DGX, NVIDIA Nsight, NVLink, NVSwitch, and Tesla 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

© 2013-2022 NVIDIA Corporation. All rights reserved.