



NVIDIA HPC SDK RELEASE NOTES

RN-09976-001-V21.3 | April 2021



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.....	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 21.3 release 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.

Key features that are new in the 21.3 release of the NVIDIA HPC SDK for Linux include:

- ▶ The HPC SDK now includes the HPC-X toolkit, a comprehensive software package that includes MPI, SHMEM/PGAS communications libraries and various acceleration packages.
- ▶ C++ Standard Language Parallelism is now supported on multicore CPUs via the `-stdpar=multicore` command line flag.
- ▶ The HPC SDK now includes 64-bit versions of the host-only OpenBLAS libraries, which have `'_ilp64.*'` suffixes on the filenames.

Chapter 2.

RELEASE COMPONENT VERSIONS

The NVIDIA HPC SDK 21.3 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 10.2	CUDA 11.0	CUDA 11.2	CUDA 10.2	CUDA 11.0	CUDA 11.2
nvc++	21.3			21.3			21.3		
nvc	21.3			21.3			21.3		
nvfortran	21.3			21.3			21.3		
nvcc	10.2.89	11.0.221	11.2.1.007	10.2.89	11.0.221	11.2.1.007	N/A	11.0.221	11.2.1.007
NCCL	2.8.3-1	2.8.3-1	2.8.3-1	2.8.3-1	2.8.3-1	2.8.3-1	N/A	N/A	2.8.3-1
NVSHMEM	1.1.3	1.1.3	1.1.3	1.1.3	1.1.3	1.1.3	N/A	N/A	N/A
cuBLAS	10.2.2.89	11.2.0.252	11.4.1.1026	10.2.2.89	11.2.0.252	11.4.1.1026	N/A	11.2.0.252	11.4.1.1026
cuFFT	10.1.2.89	10.2.1.245	10.4.0.135	10.1.2.89	10.2.1.245	10.4.0.135	N/A	10.2.1.245	10.4.0.135
cuRAND	10.1.2.89	10.2.1.245	10.2.3.135	10.1.2.89	10.2.1.245	10.2.3.135	N/A	10.2.1.245	10.2.3.135
cuSOLVER	10.3.0.89	10.6.0.245	11.1.0.135	10.3.0.89	10.6.0.245	11.1.0.135	N/A	10.6.0.245	11.1.0.135
cuSPARSE	10.3.1.89	11.1.1.245	11.4.0.135	10.3.1.89	11.1.1.245	11.4.0.135	N/A	11.1.1.245	11.4.0.135
cuTENSORFLOW	1.2.1	1.2.1	1.2.2	1.2.1	1.2.1	1.2.2	N/A	1.2.1	1.2.2
Nsight Compute	2020.3.0.0-29307467			2020.3.0.0-29307467			2020.3.1.0-29471205		
Nsight Systems	2020.5.1.85			2020.5.1.85			2021.1.1.66		
OpenMPI	3.1.5			3.1.5			3.1.5		
HPC-X	N/A	2.7.4	2.7.4	N/A	2.7.4	2.7.4	N/A	2.7.4	2.7.4
UCX	N/A	1.9.0	1.9.0	N/A	1.9.0	1.9.0	N/A	1.9.0	1.9.0
OpenBLAS	0.3.13			0.3.13			0.3.13		

	Linux_x86_64			Linux_ppc64le			Linux_aarch64		
	CUDA 10.2	CUDA 11.0	CUDA 11.2	CUDA 10.2	CUDA 11.0	CUDA 11.2	CUDA 10.2	CUDA 11.0	CUDA 11.2
Scalapack	2.1.0			2.1.0			2.1.0		
Thrust	1.9.7	1.9.9	1.10.0	1.9.7	1.9.9	1.10.0	1.9.7	1.9.10	1.10.0
CUB	N/A	1.9.9	1.10.0	N/A	1.9.9	1.10.0	N/A	1.9.9	1.10.0
libcu++	1.0.0	2.0.0	2.0.0	1.0.0	2.0.0	2.0.0	1.0.0	2.0.0	2.0.0

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 16.04, 18.04, 19.10, 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 16.04, 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	RHEL 8.1	C99: 4.8	450.36

Architecture	Linux Distributions	Minimum gcc/ glibc Toolchain	Minimum CUDA Driver
	Ubuntu 18.04	C11: 4.9 C++03: 4.8 C++11: 4.9 C++14: 5.1 C++17: 7.1	

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 21.3 includes the following CUDA toolchain versions:

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

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

Chapter 4. KNOWN LIMITATIONS

- ▶ When invoking MPI executables directly from `comm_libs/mpi/bin`, the user must explicitly set the `LD_LIBRARY_PATH` environment variable to include the `compilers/lib` directory, or use the included modulefiles with the NVIDIA HPC SDK to access the MPI executables.
- ▶ When compiling for Skylake Core or Skylake Xeon processors, it is recommended that the target processor compiler option `-tp skylake` be specified. Specifying `-tp skylake` at link time is optional.
- ▶ The `cuda-gdb` debugger is included in this release. Currently, Fortran arrays with non-constant bounds are 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).

Chapter 5.

DEPRECATIONS

- ▶ The 21.3 release of the NVIDIA HPC SDK supports NVIDIA GPUs with compute capability 3.5 (cc35) and newer.
- ▶ Support for the KNL architecture of multicore CPUs in the NVIDIA HPC SDK is deprecated in the 21.3 release.
- ▶ Starting with the 21.3 release of the NVIDIA HPC SDK, the `-cuda` option for NVC++ and NVFORTRAN will no longer automatically link the NVIDIA GPU math libraries.
- ▶ Support for the Kepler architecture of NVIDIA GPUs is deprecated starting with the 21.3 release of the NVIDIA HPC SDK.

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

© 2013-2021 NVIDIA Corporation. All rights reserved.