

## NVIDIA HPC SDK RELEASE NOTES

RN-09976-001-V20.11 | November 2020



### 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 Toolkit 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 20.11 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 20.11 release of the NVIDIA HPC SDK for Linux include:

- Beta support of a subset of OpenMP target offload features for NVIDIA GPUs on x86-based systems can be enabled in NVC++, NVC, and NVFORTRAN with -mp=gpu.
- NVFORTRAN accelerates the DO CONCURRENT construct to NVIDIA GPUs by default when enabled via the -stdpar option.
- Use the nvlamath Fortran library to access GPU-accelerated BLAS routines using the standard BLAS interfaces.
- The NVIDIA HPC SDK now supports version 11.1 of the CUDA toolchain.
- HPC SDK library and developer tool components have been updated to their newest versions.
- RPM packages for the Suse Linux Enterprise Server operating system are now available as download options.
- The default C++ Standard version for NVC++ is C++11.
- A number of new command-line options have been implemented for HPC SDK compilers to increase compatibility with GCC:
  - -f[-no]exceptions Enable[/disable] exception handling
  - -x Specify the language for any following input files, instead of letting the compiler choose based on suffix. Turn off with -x none.
  - -Wshadow Issue a warning whenever a variable hides another variable
  - -Wstrict-prototypes C only. Warn if a function is declared or defined without specifying the argument types
  - -Wno-long-long Suppress warnings about long long being nonstandard
  - -Wundef Warn about using uninitialized variables
  - -fpie Generates position independent executable. Implies -pie

# Chapter 2. RELEASE COMPONENT VERSIONS

The NVIDIA HPC SDK 20.11 release contains the following versions of each component:

	Linux_x86_64		Linux_ppc64le			Linux_aarch64			
	CUDA 10.2	CUDA 11.0	CUDA 11.1	CUDA 10.2	CUDA 11.0	CUDA 11.1	CUDA 10.2	CUDA 11.0	CUDA 11.1
nvc++	20.11		20.11			20.11			
nvc		20.11		20.11			20.11		
nvfortran		20.11		20.11		20.11			
nvcc	10.2.89	11.0.228	11.1.74	10.2.89	11.0.228	11.1.74	N/A	11.0.228	11.1.74
NCCL	2.7.8-1	2.7.8-1	2.7.8-1	2.7.8-1	2.7.8-1	2.7.8-1	N/A	N/A	2.7.8-1
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	11.2.1.74	10.2.2.89	11.2.0.252	11.2.1.74	N/A	11.2.0.252	11.2.1.74
cuFFT	10.1.2.89	10.2.1.245	10.3.0.74	10.1.2.89	10.2.1.245	10.3.0.74	N/A	10.2.1.245	10.3.0.74
cuRAND	10.1.2.89	10.2.1.245	10.2.2.74	10.1.2.89	10.2.1.245	10.2.2.74	N/A	10.2.1.245	10.2.2.74
cuSOLVER	10.3.0.89	10.6.0.245	11.0.0.74	10.3.0.89	10.6.0.245	11.0.0.74	N/A	10.6.0.245	11.0.0.74
cuSPARSE	10.3.1.89	11.1.1.245	11.2.0.275	10.3.1.89	11.1.1.245	11.2.0.275	N/A	11.1.1.245	11.2.0.275
cuTENSOF	1.2.1	1.2.1	1.2.1	1.2.1	1.2.1	1.2.1	N/A	1.2.1	1.2.1
Nsight Compute	2020.1.0.33-28294165		2020.1.0.33-28294165			2020.1.0.33-28294165			
Nsight Systems	2020.3.1.54 CLI and GUI		2020.3.1.54 CLI			2020.3.1.54 CLI			
OpenMPI	3.1.5		3.1.5		3.1.5				
OpenMPI 4 (NVHPC beta)	N/A	N/A	4.0.5	N/A	N/A	4.0.5	N/A	N/A	4.0.5
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

### Table 1 HPC SDK Release Components

	Linux_x86_64		Linux_ppc64le			Linux_aarch64			
	CUDA 10.2	CUDA 11.0	CUDA 11.1	CUDA 10.2	CUDA 11.0	CUDA 11.1	CUDA 10.2	CUDA 11.0	CUDA 11.1
OpenBLAS		0.3.10			0.3.10			0.3.10	
Scalapack		2.1.0			2.1.0			2.1.0	
Thrust	1.9.7	1.9.9	1.9.9	1.9.7	1.9.9	1.9.9	1.9.7	1.9.10	1.9.10
CUB	N/A	1.9.9	1.9.9	N/A	1.9.9	1.9.9	N/A	1.9.9	1.9.9
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

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 8.0, 8.1 Fedora 29, 30, 31, 32 OpenSUSE Leap 15.0, 15.1 RHEL 6.4, 6.5, 6.6, 6.7, 6.8, 6.9, 6.10 RHEL 7.0, 7.1, 7.2, 7.3, 7.4, 7.5, 7.6, 7.7, 7.8 RHEL 8.0, 8.1 SLES 12SP4, 12SP5, 15SP1 Ubuntu 14.04, 16.04, 18.04, 19.04, 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	418.39
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	410.45

### Table 2 HPC SDK Platform Requirements

Architecture	Linux Distributions	Minimum gcc/ glibc Toolchain	Minimum CUDA Driver
aarch64	RHEL 8.1 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	450.36

## 3.2. Supported CUDA Toolkit 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 as its first line of output. You can use it to find out which version of the CUDA Driver is installed on your system.

The NVIDIA HPC SDK 20.11 includes stand-alone support for the following CUDA toolchain versions:

- CUDA 10.2
- ► CUDA 11.0
- ► CUDA 11.1

See the NVIDIA HPC Compilers User Guide for information about using the HPC SDK Fortran, C++ and C compilers with alternative versions of the CUDA toolchain.

## Chapter 4. KNOWN LIMITATIONS

- A known issue with CUDA 11.1, GNU toolchain versions 4.x, and the NVHPC compiler option -gpu=time may cause applications built with -gpu=time to crash at runtime.
- When compiling for Skylake Core or Skylake Xeon processors, it is recommended that the target processor compiler option -tp skylake is 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).
- The Open MPI 4.0.5 build with UCX 1.9.0 is a beta feature in the NVIDIA HPC SDK 20.11. It is built with support for MOFED 5.0; users on systems using earlier MOFED versions may encounter compatibility issues.

# Chapter 5. DEPRECATIONS

- Support for the CUDA 10.1 toolchain has been dropped from the NVIDIA HPC SDK version 20.11; CUDA toolchains 10.2, 11.0 and 11.1 are now supported.
- Support for the KNL architecture of multicore CPUs in the NVIDIA HPC SDK will be deprecated in the upcoming 21.1 release.
- GCC 4.8 is the minimum required version for the NVIDIA HPC SDK 20.11.
- The NVIDIA HPC SDK will discontinue support for the Ubuntu operating system on ppc64le architecture systems starting with the upcoming 21.1 release.

### 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

© 2013-2020 NVIDIA Corporation. All rights reserved.

