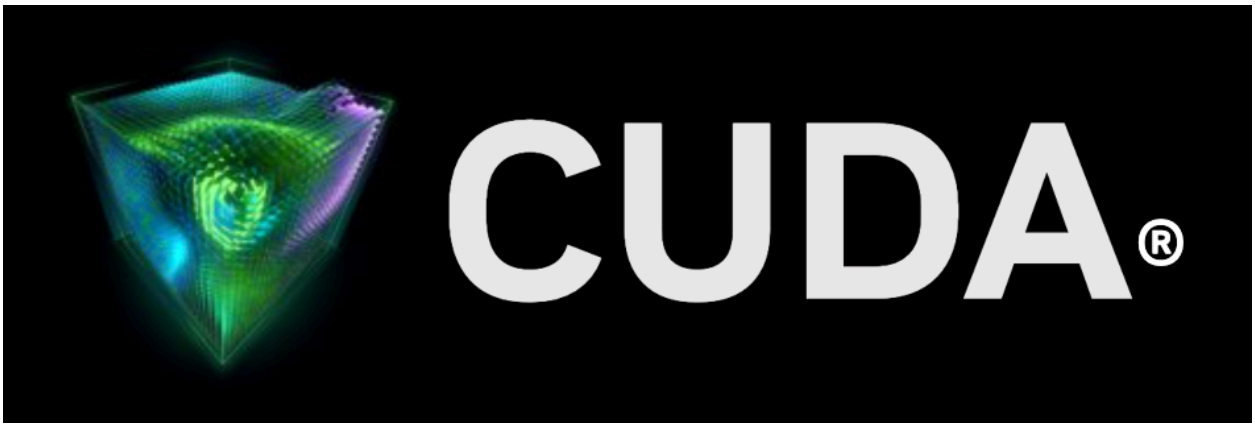

Release Notes

Release 13.4

NVIDIA Corporation

Jul 16, 2026

Contents

Chapter 1. Overview

CUDA Toolkit 13.4 Developer Preview Release Notes

CUDA Toolkit 13.4 Developer Preview provides access to the CUDA 13.4 user-space software stack before General Availability.

It is intended for developers to evaluate new APIs and PTX ISA features, porting source code, validating build systems and CI/CD workflows, updating compilers or libraries, and preparing support for newly enabled platforms.

Key areas include:

- CUDA compiler and PTX ISA updates.
- CUDA C++ and CUDA Tile programming enhancements.
- CUDA Runtime and Driver API updates.
- CUDA libraries and developer tools updates.
- CUDA Toolkit preview for RTX Spark and Rubin.

CUDA drivers, platform firmware, and system images are not included. Runtime validation may require compatible hardware and compatible drivers obtained separately.

Important: This Developer Preview is pre-release software. It may be incomplete, may contain errors or design flaws, and may differ from the final CUDA Toolkit 13.4 GA release. Use it at your own risk. Do not use it for production deployment, certification, benchmarking, or business-critical systems.

Any performance data collected with this Developer Preview is non-representative and must not be published, reported, compared, or used to characterize NVIDIA hardware and/or software.

Chapter 2. Release Highlights

2.1. CUDA Compiler and PTX

CUDA Toolkit 13.4 preview includes compiler and device-code generation updates:

- Updated host compiler and standard-library compatibility.
- CUDA C++ language support, diagnostics, and code-generation improvements.
- New and updated PTX ISA, including some Rubin capabilities.
- CUDA Tile and TileIR compilation and inspection.

2.2. CUDA C++ and CUDA Tile

CUDA 13.4 adds or updates programming capabilities related to:

- Numeric formats and mixed-precision computation.
- Matrix multiply-accumulate operations.
- Tile views, layouts, and sub-views.
- Reductions, atomic reductions, and user-defined reductions.
- Tile-level synchronization and communication.
- Compiler diagnostics and code-generation correctness.

2.3. Runtime APIs and Libraries

The preview includes the CUDA 13.4 headers, import libraries, runtime components, API definitions, Math Libraries, Domain Libraries required for compilation, and linking workflows.

2.4. Developer Tools

The preview includes updates to developer tools for compiling, linking, and inspecting CUDA applications.

2.5. RTX Spark - Windows on Arm Preview

The preview includes compiler and build-system enablement for Windows Arm64 targets.

- Native Windows Arm64 development for the upcoming RTX Spark device.
- Cross-compiling Windows Arm64 CUDA applications using the Windows x86_64 Toolkit.

For more information, see the RTX Spark Developer Preview landing page.

Chapter 3. CUDA Platform

3.1. CUDA Toolkit Major Components

Note: Starting with CUDA 11, individual components within the CUDA Toolkit (for example: compiler, libraries, tools) are versioned independently.

For CUDA 13.4, the table below indicates the versions:

Table 1: CUDA 13.4 Component Versions

Component Name		Version Information	Supported Architectures	Supported Platforms
CUDA C++ Core Compute Libraries	Thrust	3.3.3	x86_64, arm64-sbsa	Linux, Windows
	CUB	3.3.3		
	libcu++	3.3.3		
	Cooperative Groups	13.3.3.4.1		
CUDA Application Compiler (crt)		13.4.46	x86_64, arm64-sbsa	Linux, Windows
CUDA Compilation Optimizer (ctadvisor)		13.4.46	x86_64, arm64-sbsa	Linux, Windows
CUDA Runtime (cudart)		13.4.46	x86_64, arm64-sbsa	Linux, Windows
CUDA culibos		13.4.46	x86_64, arm64-sbsa	Linux
CUDA cuobjdump		13.4.46	x86_64, arm64-sbsa	Linux, Windows
CUPTI		13.4.46	x86_64, arm64-sbsa	Linux, Windows
CUDA cuxxfilt (demangler)		13.4.46	x86_64, arm64-sbsa	Linux, Windows

continues on next page

Table 1 – continued from previous page

Component Name	Version Information	Supported Architectures	Supported Platforms
CUDA Documentation	13.4.46	x86_64, arm64-sbsa	Linux, Windows
CUDA GDB	13.4.46	x86_64, arm64-sbsa	Linux
CUDA NVCC	13.4.46	x86_64, arm64-sbsa	Linux, Windows
CUDA nvdisasm	13.4.46	x86_64, arm64-sbsa	Linux, Windows
CUDA NVML Headers	13.4.46	x86_64, arm64-sbsa	Linux, Windows
CUDA nvprune	13.4.46	x86_64, arm64-sbsa	Linux, Windows
CUDA NVRTC	13.4.46	x86_64, arm64-sbsa	Linux, Windows
CUDA NVTX	13.4.46	x86_64, arm64-sbsa	Linux, Windows
CUDA OpenCL	13.4.46	x86_64	Linux, Windows
CUDA Profiler API	13.4.46	x86_64, arm64-sbsa	Linux, Windows
CUDA Sandbox dev	13.4.46	x86_64, arm64-sbsa	Linux
CUDA Compute Sanitizer API	13.4.46	x86_64, arm64-sbsa	Linux, Windows
CUDA TILE-IR AS	13.4.46	x86_64, arm64-sbsa	Linux, Windows
CUDA cuBLAS	13.7.0.10	x86_64, arm64-sbsa	Linux, Windows
CUDA cuDLA	13.4.46	x86_64, arm64-sbsa	Linux
CUDA cuFFT	12.4.0.16	x86_64, arm64-sbsa	Linux, Windows
CUDA cuFile	1.19.0.108	x86_64, arm64-sbsa	Linux
CUDA cuobjclient	1.3.0.108	x86_64, arm64-sbsa	Linux
CUDA cuRAND	10.4.4.46	x86_64, arm64-sbsa	Linux, Windows

continues on next page

Table 1 – continued from previous page

Component Name	Version Information	Supported Architectures	Supported Platforms
CUDA cuSOLVER	12.3.1.12	x86_64, arm64-sbsa	Linux, Windows
CUDA cuSPARSE	12.8.6.46	x86_64, arm64-sbsa	Linux, Windows
CUDA NPP	13.2.0.32	x86_64, arm64-sbsa	Linux, Windows
CUDA nvFatbin	13.4.46	x86_64, arm64-sbsa	Linux, Windows
CUDA nvJitLink	13.4.46	x86_64, arm64-sbsa	Linux, Windows
CUDA nvJPEG	13.2.2.32	x86_64, arm64-sbsa	Linux, Windows
CUDA nvptxcompiler	13.4.46	x86_64, arm64-sbsa	Linux, Windows
CUDA nvvm	13.4.46	x86_64, arm64-sbsa	Linux, Windows
Nsight Compute	2026.3.0.11	x86_64, arm64-sbsa	Linux, Windows
Nsight Systems	2026.3.2.239	x86_64, arm64-sbsa	Linux, Windows
Nsight Visual Studio Edition (VSE)	2026.3.0.26163	x86_64 (Windows)	Windows
nvidia_fs[1]	2.29.4	x86_64, arm64-sbsa	Linux
nvlsm	2025.10.12	x86_64, arm64-sbsa	Linux
Visual Studio Integration	13.4.46	x86_64 (Windows)	Windows

3.2. CUDA Driver

Note:

- Starting with CUDA 13.4, the NVIDIA Linux driver is no longer bundled with the CUDA Toolkit.
- Download and install it separately from the official NVIDIA driver download page.

[1] Only available on select Linux distros

Running a CUDA application requires the system with at least one CUDA-capable GPU and a driver that is compatible with the CUDA Toolkit. See Table 3. For more information about various GPU products that are CUDA-capable, visit https://developer.nvidia.com/cuda-gpus.

Each release of the CUDA Toolkit requires a minimum version of the CUDA driver. The CUDA driver is backward compatible, meaning that applications compiled against a particular CUDA Toolkit version will continue to work on subsequent (later) driver releases.

More information on compatibility can be found at https://docs.nvidia.com/cuda/cuda-c-best-practices-guide/index.html#cuda-compatibility-and-upgrades.

Note: Starting with CUDA 11.0, the toolkit components are individually versioned, and the toolkit itself is versioned as shown in the table below.

The minimum required driver version for CUDA minor version compatibility is shown below. CUDA minor version compatibility is described in detail in https://docs.nvidia.com/deploy/cuda-compatibility/index.html

Table 2: CUDA Toolkit and Minimum Required Driver Version for CUDA Minor Version Compatibility

CTK Version	Driver Range for Minor Version Compatibility	
	Min	**Max**
13.x	>= 580	N/A
12.x	>= 525	< 580
11.x	>= 450	< 525

Note: Driver version 580 or later supports CUDA 13.x minor-version compatibility. CUDA 13.4 features and newly enabled platforms require an R616 or later CUDA 13.4-compatible driver.

* Using a Minimum Required Version that is **different** from Corresponding Driver Version could be allowed in compatibility mode – please read the CUDA Compatibility Guide for details.

** Starting with CUDA 13.1, the NVIDIA GPU driver on Windows is **no longer** bundled with the CUDA Toolkit package. Users must download and install the appropriate NVIDIA driver separately from the official driver download page.

For more information on supported driver versions, see the CUDA Compatibility Guide for drivers.

*** CUDA 11.0 shipped with earlier driver versions. Minor-version compatibility across the CUDA 11.x family requires driver version 450.80.02 or later on Linux, or 452.39 or later on Windows.

The version of the NVIDIA GPU driver associated with each CUDA Toolkit release is shown below.

Table 3: CUDA Toolkit and Corresponding Driver Versions

CUDA Toolkit	Corresponding Driver Version	
	Linux x86_64 Driver Version	Windows x86_64 Driver Version
CUDA 13.4 Developer Preview	N/A[2]	N/A
CUDA 13.3 Update 1	>=610.43.02	N/A

continues on next page

Table 3 – continued from previous page

CUDA Toolkit	Corresponding Driver Version	
CUDA 13.3 GA	>=610.43.02	N/A
CUDA 13.2 Update 1	>=595.58.03	N/A
CUDA 13.2 GA	>=595.45.04	N/A
CUDA 13.1 Update 1	>=590.48.01	N/A
CUDA 13.1 GA	>=590.44.01	N/A
CUDA 13.0 Update 2	>=580.95.05	N/A
CUDA 13.0 Update 1	>=580.82.07	N/A
CUDA 13.0 GA	>=580.65.06	N/A

Older CUDA versions (12.9 and earlier)

Table 4: Corresponding Driver Versions (CUDA 12.9 and earlier)

CUDA Toolkit	Corresponding Driver Version	
	Linux x86_64 Driver Version	Windows x86_64 Driver Version
CUDA 12.9 Update 1	>=575.57.08	>=576.57
CUDA 12.9 GA	>=575.51.03	>=576.02
CUDA 12.8 Update 1	>=570.124.06	>=572.61
CUDA 12.8 GA	>=570.26	>=570.65
CUDA 12.6 Update 3	>=560.35.05	>=561.17
CUDA 12.6 Update 2	>=560.35.03	>=560.94
CUDA 12.6 Update 1	>=560.35.03	>=560.94
CUDA 12.6 GA	>=560.28.03	>=560.76
CUDA 12.5 Update 1	>=555.42.06	>=555.85
CUDA 12.5 GA	>=555.42.02	>=555.85
CUDA 12.4 Update 1	>=550.54.15	>=551.78
CUDA 12.4 GA	>=550.54.14	>=551.61
CUDA 12.3 Update 1	>=545.23.08	>=546.12
CUDA 12.3 GA	>=545.23.06	>=545.84
CUDA 12.2 Update 2	>=535.104.05	>=537.13

continues on next page

[2] The CUDA 13.4 driver is not bundled with the Toolkit. Install the separately-released NVIDIA Developer Driver (R616, >=616.00). Existing CUDA 13.x applications run on drivers >=580 under CUDA compatibility; new 13.4 features and newly enabled platforms require >=616.00.

Table 4 – continued from previous page

CUDA Toolkit	Corresponding Driver Version	
CUDA 12.2 Update 1	>=535.86.09	>=536.67
CUDA 12.2 GA	>=535.54.03	>=536.25
CUDA 12.1 Update 1	>=530.30.02	>=531.14
CUDA 12.1 GA	>=530.30.02	>=531.14
CUDA 12.0 Update 1	>=525.85.12	>=528.33
CUDA 12.0 GA	>=525.60.13	>=527.41
CUDA 11.8 GA	>=520.61.05	>=520.06
CUDA 11.7 Update 1	>=515.48.07	>=516.31
CUDA 11.7 GA	>=515.43.04	>=516.01
CUDA 11.6 Update 2	>=510.47.03	>=511.65
CUDA 11.6 Update 1	>=510.47.03	>=511.65
CUDA 11.6 GA	>=510.39.01	>=511.23
CUDA 11.5 Update 2	>=495.29.05	>=496.13
CUDA 11.5 Update 1	>=495.29.05	>=496.13
CUDA 11.5 GA	>=495.29.05	>=496.04
CUDA 11.4 Update 4	>=470.82.01	>=472.50
CUDA 11.4 Update 3	>=470.82.01	>=472.50
CUDA 11.4 Update 2	>=470.57.02	>=471.41
CUDA 11.4 Update 1	>=470.57.02	>=471.41
CUDA 11.4.0 GA	>=470.42.01	>=471.11
CUDA 11.3.1 Update 1	>=465.19.01	>=465.89
CUDA 11.3.0 GA	>=465.19.01	>=465.89
CUDA 11.2.2 Update 2	>=460.32.03	>=461.33
CUDA 11.2.1 Update 1	>=460.32.03	>=461.09
CUDA 11.2.0 GA	>=460.27.03	>=460.82
CUDA 11.1.1 Update 1	>=455.32	>=456.81
CUDA 11.1 GA	>=455.23	>=456.38
CUDA 11.0.3 Update 1	>= 450.51.06	>= 451.82
CUDA 11.0.2 GA	>= 450.51.05	>= 451.48
CUDA 11.0.1 RC	>= 450.36.06	>= 451.22
CUDA 10.2.89	>= 440.33	>= 441.22

continues on next page

Table 4 – continued from previous page

CUDA Toolkit	Corresponding Driver Version	
CUDA 10.1 (10.1.105 general release, and updates)	>= 418.39	>= 418.96
CUDA 10.0.130	>= 410.48	>= 411.31
CUDA 9.2 (9.2.148 Update 1)	>= 396.37	>= 398.26
CUDA 9.2 (9.2.88)	>= 396.26	>= 397.44
CUDA 9.1 (9.1.85)	>= 390.46	>= 391.29
CUDA 9.0 (9.0.76)	>= 384.81	>= 385.54
CUDA 8.0 (8.0.61 GA2)	>= 375.26	>= 376.51
CUDA 8.0 (8.0.44)	>= 367.48	>= 369.30
CUDA 7.5 (7.5.16)	>= 352.31	>= 353.66
CUDA 7.0 (7.0.28)	>= 346.46	>= 347.62

- CUDA Toolkit driver bundling (pre-CUDA 13.1):
  - The CUDA Toolkit previously included an NVIDIA GPU driver for convenience.
  - This bundled driver was intended only for development purposes.
  - It is not recommended for production use, especially with NVIDIA data center GPUs.
- Recommended driver for NVIDIA data center GPUs:
  - For production environments using NVIDIA data center GPUs, download the latest production driver from the official NVIDIA Driver Downloads page:

    https://www.nvidia.com/drivers
- Optional driver installation during Toolkit setup:
  - During CUDA Toolkit installation, users may choose to skip driver installation:
    - On Windows: via interactive or silent install options.
    - On Linux: by skipping driver meta packages.
- Change in CUDA 13.1 (Windows-specific):
  - Starting with CUDA 13.1, the NVIDIA GPU driver on Windows is **no longer bundled** with the CUDA Toolkit.
  - Windows users must **manually download and install** the appropriate driver from the official NVIDIA site.
- Change in CUDA 13.4 (Linux-specific):
  - Starting with CUDA 13.4, the NVIDIA Linux driver is released separately as the **NVIDIA Developer Driver** (R616), independent of the CUDA Toolkit.
  - Linux users must **manually download and install** it from the official NVIDIA driver download page.
- Driver compatibility notes:

- Some compatibility tables list “N/A” for a driver version – for Windows, and for Linux in the CUDA 13.4 Developer Preview, where the driver ships separately as the NVIDIA Developer Driver (R616).
- Users must still ensure the installed driver meets or exceeds the minimum required version for the CUDA Toolkit.
- For details, refer to the official CUDA Compatibility Guide for Drivers:

  https://docs.nvidia.com/deploy/cuda-compatibility/index.html

Chapter 4. Packaging and Distribution

CUDA Toolkit 13.4 Developer Preview is distributed through Preview-labeled package locations on packages.nvidia.com.

4.1. Install Packages

Table 5: CUDA Toolkit 13.4 Developer Preview Install Packages

Installer type	**Packages**
Linux DEB repository	Network repository for x86_64, ARM64 SBSA
Linux RPM repository	Network repository for x86_64, ARM64 SBSA
Windows installer	Local installer for Windows x86_64, Windows ARM64
Binary archives	Linux x86_64, Linux ARM64 SBSA, Windows x86_64, Windows ARM64

Chapter 5. Driver Compatibility

For this Developer Preview:

- The CUDA driver is not included with the Developer Preview.
- Existing CUDA functionality runs on compatible existing drivers under applicable CUDA compatibility rules.
- New CUDA 13.4 features and newly enabled platforms require a CUDA 13.4-compatible driver.

Chapter 6. Known Limitations

- This preview has not completed full validation. It may be incomplete, may contain errors or design flaws.
- APIs, PTX ISA features, compiler behavior, package contents, and documentation may change before GA.
- Platform, library, algorithm, data-type, and developer-tool coverage may be incomplete.
- Architecture-specific optimizations are incomplete or disabled.
- Unsupported component or platform combinations may fail during compilation, linking, loading, or execution.
- Profiling, debugging, tracing, and sanitization coverage are limited.

Chapter 7. Notices

7.1. Notice

This document is provided for information purposes only and shall not be regarded as a warranty of a certain functionality, condition, or quality of a product. NVIDIA Corporation ("NVIDIA") makes no representations or warranties, expressed or implied, as to the accuracy or completeness of the information contained in this document and assumes no responsibility for any errors contained herein. NVIDIA shall have no liability for the consequences or use of such information or for any infringement of patents or other rights of third parties that may result from its use. This document is not a commitment to develop, release, or deliver any Material (defined below), code, or functionality.

NVIDIA reserves the right to make corrections, modifications, enhancements, improvements, and any other changes to this document, at any time without notice.

Customer should obtain the latest relevant information before placing orders and should verify that such information is current and complete.

NVIDIA products are sold subject to the NVIDIA standard terms and conditions of sale supplied at the time of order acknowledgement, unless otherwise agreed in an individual sales agreement signed by authorized representatives of NVIDIA and customer ("Terms of Sale"). NVIDIA hereby expressly objects to applying any customer general terms and conditions with regards to the purchase of the NVIDIA product referenced in this document. No contractual obligations are formed either directly or indirectly by this document.

NVIDIA products are not designed, authorized, or warranted to be suitable for use in medical, military, aircraft, space, or life support equipment, nor in applications where failure or malfunction of the NVIDIA product can reasonably be expected to result in personal injury, death, or property or environmental damage. NVIDIA accepts no liability for inclusion and/or use of NVIDIA products in such equipment or applications and therefore such inclusion and/or use is at customer's own risk.

NVIDIA makes no representation or warranty that products based on this document will be suitable for any specified use. Testing of all parameters of each product is not necessarily performed by NVIDIA. It is customer's sole responsibility to evaluate and determine the applicability of any information contained in this document, ensure the product is suitable and fit for the application planned by customer, and perform the necessary testing for the application in order to avoid a default of the application or the product. Weaknesses in customer's product designs may affect the quality and reliability of the NVIDIA product and may result in additional or different conditions and/or requirements beyond those contained in this document. NVIDIA accepts no liability related to any default, damage, costs, or problem which may be based on or attributable to: (i) the use of the NVIDIA product in any manner that is contrary to this document or (ii) customer product designs.

No license, either expressed or implied, is granted under any NVIDIA patent right, copyright, or other NVIDIA intellectual property right under this document. Information published by NVIDIA regarding third-party products or services does not constitute a license from NVIDIA to use such products or

services or a warranty or endorsement thereof. Use of such information may require a license from a third party under the patents or other intellectual property rights of the third party, or a license from NVIDIA under the patents or other intellectual property rights of NVIDIA.

Reproduction of information in this document is permissible only if approved in advance by NVIDIA in writing, reproduced without alteration and in full compliance with all applicable export laws and regulations, and accompanied by all associated conditions, limitations, and notices.

THIS DOCUMENT AND 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. TO THE EXTENT NOT PROHIBITED BY LAW, IN NO EVENT WILL NVIDIA BE LIABLE FOR ANY DAMAGES, INCLUDING WITHOUT LIMITATION ANY DIRECT, INDIRECT, SPECIAL, INCIDENTAL, PUNITIVE, OR CONSEQUENTIAL DAMAGES, HOWEVER CAUSED AND REGARDLESS OF THE THEORY OF LIABILITY, ARISING OUT OF ANY USE OF THIS DOCUMENT, EVEN IF NVIDIA HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. Notwithstanding any damages that customer might incur for any reason whatsoever, NVIDIA's aggregate and cumulative liability towards customer for the products described herein shall be limited in accordance with the Terms of Sale for the product.

7.2. OpenCL

OpenCL is a trademark of Apple Inc. used under license to the Khronos Group Inc.

7.3. Trademarks

NVIDIA and the NVIDIA logo are trademarks 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