



# cuDNN Support Matrix

# Table of Contents

<b>Chapter 1. Hardware Requirements.....</b>	<b>1</b>
1.1. cuDNN 8.0.5.....	1
1.2. cuDNN 8.0.2 - 8.0.3.....	2
1.3. cuDNN 8.0.0 - 8.0.1 Preview.....	3
1.4. cuDNN 7.6.4 - 7.6.5.....	3
1.5. cuDNN 7.6.3.....	5
1.6. cuDNN 7.5.1 - 7.6.2.....	6
1.7. cuDNN 7.3.1 - 7.5.0.....	7
1.8. cuDNN 7.1.4 - 7.2.1.....	8
<b>Chapter 2. Software Requirements.....</b>	<b>9</b>
2.1. Linux.....	9
2.1.1. cuDNN 8.0.5 For Linux.....	9
2.1.2. cuDNN 8.0.4 For Linux.....	9
2.1.3. cuDNN 8.0.2 - 8.0.3 For Linux.....	10
2.1.4. cuDNN 8.0.0 - 8.0.1 Preview For Linux.....	10
2.1.5. cuDNN 7.6.4 - 7.6.5 For Linux.....	11
2.2. Windows.....	11
2.2.1. cuDNN 8.x.x For Windows.....	11
2.2.2. cuDNN 7.6.4 - 7.6.5 For Windows.....	12

---

# Chapter 1. Hardware Requirements

The following sections highlight the compatibility of cuDNN versions with the various supported CUDA, CUDA driver, and NVIDIA™ hardware versions.

## 1.1. cuDNN 8.0.5

Supported NVIDIA hardware, CUDA, and CUDA driver versions for the cuDNN 8.0.44 and 8.0.5 releases.

Table 1. Supported NVIDIA hardware, CUDA, and CUDA driver

Supported NVIDIA Hardware	CUDA Version	CUDA Compute Capability	CUDA Driver Version
<ul style="list-style-type: none"><li>▶ NVIDIA Ampere GPU architecture</li><li>▶ Turing</li><li>▶ Volta</li><li>▶ Pascal</li><li>▶ Maxwell</li><li>▶ Kepler</li></ul>	CUDA 11.1	SM 3.5 and later	r450, r455
<ul style="list-style-type: none"><li>▶ NVIDIA Ampere GPU architecture</li><li>▶ Turing</li><li>▶ Volta</li><li>▶ Pascal</li><li>▶ Maxwell</li><li>▶ Kepler</li></ul>	CUDA 11.0	SM 3.5 and later	r450
<ul style="list-style-type: none"><li>▶ Turing</li><li>▶ Volta</li><li>▶ Xavier</li></ul>	CUDA 10.2	SM 3.0 and later	r440

Supported NVIDIA Hardware	CUDA Version	CUDA Compute Capability	CUDA Driver Version
<ul style="list-style-type: none"> <li>▶ Pascal</li> <li>▶ Maxwell</li> <li>▶ Kepler</li> </ul>			
<ul style="list-style-type: none"> <li>▶ Turing</li> <li>▶ Volta</li> <li>▶ Pascal</li> <li>▶ Maxwell</li> <li>▶ Kepler</li> </ul>	CUDA 10.1 Update 2	SM 3.0 and later	r418.39

## 1.2. cuDNN 8.0.2 - 8.0.3

Supported NVIDIA hardware, CUDA, and CUDA driver versions for the cuDNN version 8.0.2 and 8.0.3 releases.

Table 2. Supported NVIDIA hardware, CUDA, and CUDA driver

Supported NVIDIA Hardware	CUDA Version	CUDA Compute Capability	CUDA Driver Version
<ul style="list-style-type: none"> <li>▶ NVIDIA Ampere architecture-based A100/GA100 GPU</li> <li>▶ Turing</li> <li>▶ Volta</li> <li>▶ Pascal</li> <li>▶ Maxwell</li> <li>▶ Kepler</li> </ul>	CUDA 11.0	SM 3.5 and later	r450
<ul style="list-style-type: none"> <li>▶ Turing</li> <li>▶ Volta</li> <li>▶ Xavier</li> <li>▶ Pascal</li> <li>▶ Maxwell</li> <li>▶ Kepler</li> </ul>	CUDA 10.2	SM 3.0 and later	r440
<ul style="list-style-type: none"> <li>▶ Turing</li> <li>▶ Volta</li> </ul>	CUDA 10.1 Update 2	SM 3.0 and later	r418.39

Supported NVIDIA Hardware	CUDA Version	CUDA Compute Capability	CUDA Driver Version
<ul style="list-style-type: none"> <li>▶ Pascal</li> <li>▶ Maxwell</li> <li>▶ Kepler</li> </ul>			

## 1.3. cuDNN 8.0.0 - 8.0.1 Preview

Supported NVIDIA hardware, CUDA, and CUDA driver versions for cuDNN versions 8.0.0 and 8.0.1 Preview releases.

Table 3. Supported NVIDIA hardware, CUDA, and CUDA driver

Supported NVIDIA Hardware	CUDA Version	CUDA Compute Capability	CUDA Driver Version
<ul style="list-style-type: none"> <li>▶ NVIDIA Ampere architecture-based A100/GA100 GPU</li> <li>▶ Turing</li> <li>▶ Volta</li> <li>▶ Pascal</li> <li>▶ Maxwell</li> <li>▶ Kepler</li> </ul>	CUDA 11.0	SM 3.5 and later	r450
<ul style="list-style-type: none"> <li>▶ Turing</li> <li>▶ Volta</li> <li>▶ Xavier</li> <li>▶ Pascal</li> <li>▶ Maxwell</li> <li>▶ Kepler</li> </ul>	CUDA 10.2	SM 3.0 and later	r440

## 1.4. cuDNN 7.6.4 - 7.6.5

Refer to the following table to view the list of supported NVIDIA hardware, CUDA, and CUDA driver versions for cuDNN versions 7.6.4 and v7.6.5.

Table 4. Supported NVIDIA hardware, CUDA, and CUDA driver versions for cuDNN versions 7.6.4 and 7.6.5

Supported NVIDIA Hardware (Compute Capability)	CUDA Version	CUDA Driver Version
<ul style="list-style-type: none"> <li>▶ Turing (7.5)</li> <li>▶ Volta (7.x)</li> <li>▶ Xavier (7.2)</li> <li>▶ Pascal (6.x)</li> <li>▶ Maxwell (5.x)</li> <li>▶ Kepler (3.x)</li> </ul>	CUDA 10.2	r440
<ul style="list-style-type: none"> <li>▶ Turing (7.5)</li> <li>▶ Volta (7.x)</li> <li>▶ Xavier (7.2)</li> <li>▶ Pascal (6.x)</li> <li>▶ Maxwell (5.x)</li> <li>▶ Kepler (3.x)</li> </ul>	CUDA 10.1.243	r418.39
<ul style="list-style-type: none"> <li>▶ Turing (7.5)</li> <li>▶ Volta (7.x)</li> <li>▶ Xavier (7.2)</li> <li>▶ Pascal (6.x)</li> <li>▶ Maxwell (5.x)</li> <li>▶ Kepler (3.x)</li> </ul>	CUDA 10.0.x	r410.48
<ul style="list-style-type: none"> <li>▶ Volta (7.x)</li> <li>▶ Xavier (7.2)</li> <li>▶ Pascal (6.x)</li> <li>▶ Maxwell (5.x)</li> <li>▶ Kepler (3.x)</li> </ul>	CUDA 9.2.148	r396.26
	CUDA 9.1.85 (Not Supported)	
<ul style="list-style-type: none"> <li>▶ Volta (7.x)</li> <li>▶ Xavier (7.2)</li> <li>▶ Pascal (6.x)</li> <li>▶ Maxwell (5.x)</li> <li>▶ Kepler (3.x)</li> </ul>	CUDA 9.0.176	r384.111

Supported NVIDIA Hardware (Compute Capability)	CUDA Version	CUDA Driver Version
	CUDA 8.0.61 (Not Supported)	

## 1.5. cuDNN 7.6.3

Refer to the following table to view the list of supported NVIDIA hardware, CUDA, and CUDA driver versions for cuDNN version 7.6.3.

Table 5. Supported NVIDIA hardware, CUDA, and CUDA driver versions for cuDNN version 7.6.3

Supported NVIDIA Hardware (Compute Capability)	CUDA Version	CUDA Driver Version
<ul style="list-style-type: none"> <li>▶ Turing (7.5)</li> <li>▶ Volta (7.x)</li> <li>▶ Xavier (7.2)</li> <li>▶ Pascal (6.x)</li> <li>▶ Maxwell (5.x)</li> <li>▶ Kepler (3.x)</li> </ul>	CUDA 10.1.243	r418.39
<ul style="list-style-type: none"> <li>▶ Turing (7.5)</li> <li>▶ Volta (7.x)</li> <li>▶ Xavier (7.2)</li> <li>▶ Pascal (6.x)</li> <li>▶ Maxwell (5.x)</li> <li>▶ Kepler (3.x)</li> </ul>	CUDA 10.0.x	r410.48
<ul style="list-style-type: none"> <li>▶ Volta (7.x)</li> <li>▶ Xavier (7.2)</li> <li>▶ Pascal (6.x)</li> <li>▶ Maxwell (5.x)</li> <li>▶ Kepler (3.x)</li> </ul>	CUDA 9.2.148	r396.26
	CUDA 9.1.85 (Not Supported)	
<ul style="list-style-type: none"> <li>▶ Volta (7.x)</li> <li>▶ Xavier (7.2)</li> <li>▶ Pascal (6.x)</li> <li>▶ Maxwell (5.x)</li> </ul>	CUDA 9.0.176	r384.111

Supported NVIDIA Hardware (Compute Capability)	CUDA Version	CUDA Driver Version
▶ Kepler (3.x)		
	CUDA 8.0.61 (Not Supported)	

## 1.6. cuDNN 7.5.1 - 7.6.2

Refer to the following table to view the list of supported NVIDIA hardware, CUDA, and CUDA driver versions for cuDNN versions 7.5.1 - 7.6.2.

Table 6. Supported NVIDIA hardware, CUDA, and CUDA driver versions for cuDNN versions 7.5.1 - 7.6.2

Supported NVIDIA Hardware (Compute Capability)	CUDA Version	CUDA Driver Version
▶ Turing (7.5) ▶ Volta (7.x) ▶ Xavier (7.2) ▶ Pascal (6.x) ▶ Maxwell (5.x) ▶ Kepler (3.x)	CUDA 10.1.105	r418.39
▶ Turing (7.5) ▶ Volta (7.x) ▶ Xavier (7.2) ▶ Pascal (6.x) ▶ Maxwell (5.x) ▶ Kepler (3.x)	CUDA 10.0.x	r410.48
▶ Volta (7.x) ▶ Xavier (7.2) ▶ Pascal (6.x) ▶ Maxwell (5.x) ▶ Kepler (3.x)	CUDA 9.2.88	r396.26
	CUDA 9.1.85 (Not Supported)	
▶ Volta (7.x) ▶ Xavier (7.2)	CUDA 9.0.176	r384.111



Supported NVIDIA Hardware (Compute Capability)	CUDA Version	CUDA Driver Version
<ul style="list-style-type: none"> <li>▶ Pascal (6.x)</li> <li>▶ Maxwell (5.x)</li> <li>▶ Kepler (3.x)</li> </ul>		
	CUDA 8.0.61 (Not Supported)	

## 1.7. cuDNN 7.3.1 - 7.5.0

Refer to the following table to view the list of supported NVIDIA hardware, CUDA, and CUDA driver versions for cuDNN versions 7.3.1 - 7.5.0.

Table 7. Supported NVIDIA hardware, CUDA, and CUDA driver versions for cuDNN versions 7.3.1 - 7.5.0

Supported NVIDIA Hardware (Compute Capability)	CUDA Version	CUDA Driver Version
<ul style="list-style-type: none"> <li>▶ Turing (7.5)</li> <li>▶ Volta (7.x)</li> <li>▶ Xavier (7.2)</li> <li>▶ Pascal (6.x)</li> <li>▶ Maxwell (5.x)</li> <li>▶ Kepler (3.x)</li> </ul>	CUDA 10.0.x	r410.48
<ul style="list-style-type: none"> <li>▶ Volta (7.x)</li> <li>▶ Xavier (7.2)</li> <li>▶ Pascal (6.x)</li> <li>▶ Maxwell (5.x)</li> <li>▶ Kepler (3.x)</li> </ul>	CUDA 9.2.88	r396.26
	CUDA 9.1.85 (Not Supported)	
<ul style="list-style-type: none"> <li>▶ Volta (7.x)</li> <li>▶ Xavier (7.2)</li> <li>▶ Pascal (6.x)</li> <li>▶ Maxwell (5.x)</li> <li>▶ Kepler (3.x)</li> </ul>	CUDA 9.0.176	r384.111
	CUDA 8.0.61 (Not Supported)	

## 1.8. cuDNN 7.1.4 - 7.2.1

Refer to the following table to view the list of supported NVIDIA hardware, CUDA, and CUDA driver versions for cuDNN versions 7.1.4 - 7.2.1.

Table 8. Supported NVIDIA hardware, CUDA, and CUDA driver versions for cuDNN versions 7.1.4 - 7.2.1

Supported NVIDIA Hardware (Compute Capability)	CUDA Version	CUDA Driver Version
	CUDA 10.0.x (Not Supported)	
<ul style="list-style-type: none"> <li>▶ Volta (7.x)</li> <li>▶ Xavier (7.2)</li> <li>▶ Pascal (6.x)</li> <li>▶ Maxwell (5.x)</li> <li>▶ Kepler (3.x)</li> </ul>	CUDA 9.2.88	r396.26
	CUDA 9.1.85 (Not Supported)	
<ul style="list-style-type: none"> <li>▶ Volta (7.x)</li> <li>▶ Xavier (7.2)</li> <li>▶ Pascal (6.x)</li> <li>▶ Maxwell (5.x)</li> <li>▶ Kepler (3.x)</li> </ul>	CUDA 9.0.176	r384.111
<ul style="list-style-type: none"> <li>▶ Xavier (7.2)</li> <li>▶ Pascal (6.x)</li> <li>▶ Maxwell (5.x)</li> <li>▶ Kepler (3.x)</li> </ul>	CUDA 8.0.61	r375.88

---

# Chapter 2. Software Requirements

The following tables highlight the compatibility of cuDNN versions with the various supported OS versions.

## 2.1. Linux

Refer to the following table to view the list of supported Linux versions for cuDNN.

### 2.1.1. cuDNN 8.0.5 For Linux

Linux versions for cuDNN 8.0.5 release

Architecture	OS Name	OS Version	Distro Information		
			Kernel	GCC	Glibc
x86_64	RHEL	7.8	3.10.0	4.8.5	2.17
		8.2 <sup>1</sup>	4.18	8.3.1	2.28
	Ubuntu	20.04	5.4.0	9.3.0	2.32
		18.04.5 LTS	4.15.0	8.2.0	2.27
		16.04.6 LTS	4.5.0	5.4.0	2.23
ppc64le	RHEL	8.2	4.18	8.3.1	2.28
AArch64 SBSA	RHEL	8	4.18	8.3.0	2.28
	Ubuntu	18.04	4.15	8.3.0	2.27
AArch64 <sup>2</sup>	Ubuntu	18.04	4.15	7.3.1	2.27

### 2.1.2. cuDNN 8.0.4 For Linux

<sup>1</sup> Not supported in CUDA 10.1 Update 2.

<sup>2</sup> Supported in CUDA 11.0 Toolkit only.

## Linux versions for cuDNN 8.0.4 release

Architecture	OS Name	OS Version	Distro Information		
			Kernel	GCC	Glibc
x86_64	RHEL	7.8	3.10.0	4.8.5	2.17
		8.2 <sup>3</sup>	4.18	8.3.1	2.28
	Ubuntu	18.04.4 LTS	4.15.0	8.2.0	2.27
		16.04.6 LTS	4.5.0	5.4.0	2.23
ppc64le	RHEL	8.2	4.18	8.3.1	2.28
	Ubuntu	18.04.4 LTS	4.4.0	5.4.0	2.27
AArch64 SBSA	RHEL	8	4.18	8.3.0	2.28
	Ubuntu	18.04	4.15	8.3.0	2.27
AArch64 <sup>2</sup>	Ubuntu	18.04			

## 2.1.3. cuDNN 8.0.2 - 8.0.3 For Linux

### Linux versions for cuDNN 8.0.2 - 8.0.3 releases

Architecture	OS Name	OS Version	Distro Information		
			Kernel	GCC	Glibc
x86_64	RHEL	7.6	3.10.0	4.8.5	2.17
		8.1	4.18	8.3.1	2.28
	Ubuntu	18.04.4 LTS	4.15.0	8.2.0	2.27
		16.04.6 LTS	4.5.0	5.4.0	2.23
ppc64le	Ubuntu	18.04.4 LTS	4.4.0	5.4.0	2.27
	RHEL	7.6			
		8.1			
AArch64 <sup>2</sup>	Ubuntu18_04				

## 2.1.4. cuDNN 8.0.0 - 8.0.1 Preview For Linux

<sup>3</sup> Not supported in CUDA 10.1 Update 2.

## Linux versions for cuDNN 8.0.0 - 8.0.1 Preview releases

Architecture	OS Name	OS Version	Distro Information		
			Kernel	GCC	Glibc
x86_64	RHEL	7.6	3.10.0	4.8.5	2.17
		8.1	4.18	8.3.1	2.28
	Ubuntu	18.04.3 LTS	4.15.0	8.2.0	2.27
		16.04.6 LTS	4.5.0	5.4.0	2.23
ppc64le	Ubuntu	18.04.3 LTS	4.4.0	5.4.0	2.27
	RHEL7	7.6			
		8.1			
AArch64 <sup>4</sup>	Ubuntu18_04				

## 2.1.5. cuDNN 7.6.4 - 7.6.5 For Linux

### Linux versions for cuDNN 7.6.4 - 7.6.5 releases

Architecture	OS Name	OS Version	Distro Information		
			Kernel	GCC	Glibc
x86_64	RHEL	7.6	3.10.0	4.8.5	2.17
	Ubuntu	18.04.3 LTS	4.15.0	8.2.0	2.27
		16.04.6 LTS	4.5.0	5.4.0	2.23
ppc64le	Ubuntu	18.04.3 LTS	4.4.0	5.4.0	2.27
	RHEL	7.6			
AArch64 <sup>2</sup>	Ubuntu18_04				

## 2.2. Windows

Refer to the following tables to view the list of supported Windows versions for cuDNN.

### 2.2.1. cuDNN 8.x.x For Windows

Windows 10 and Windows Server 2019 and 2016 are supported.

<sup>4</sup> Supported in CUDA 10.2 Toolkit only.

### Visual Studio versions based on your version of CUDA

	CUDA 11.0	CUDA 10.2
Visual Studio	2015	2012

## 2.2.2. cuDNN 7.6.4 - 7.6.5 For Windows

Windows 10 and Windows Server 2016 are supported.

### Visual Studio versions based on your version of CUDA

	CUDA 10.2	CUDA 10.1 Update 2
Visual Studio	2012	

## 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 NON-INFRINGEMENT, 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.

## VESA DisplayPort

DisplayPort and DisplayPort Compliance Logo, DisplayPort Compliance Logo for Dual-mode Sources, and DisplayPort Compliance Logo for Active Cables are trademarks owned by the Video Electronics Standards Association in the United States and other countries.

## HDMI

HDMI, the HDMI logo, and High-Definition Multimedia Interface are trademarks or registered trademarks of HDMI Licensing LLC.

## ARM

ARM, AMBA and ARM Powered are registered trademarks of ARM Limited. Cortex, MPCore and Mali are trademarks of ARM Limited. All other brands or product names are the property of their respective holders. "ARM" is used to represent ARM Holdings plc; its operating company ARM Limited; and the regional subsidiaries ARM Inc.; ARM KK; ARM Korea Limited.; ARM Taiwan Limited; ARM France SAS; ARM Consulting (Shanghai) Co. Ltd.; ARM Germany GmbH; ARM Embedded Technologies Pvt. Ltd.; ARM Norway, AS and ARM Sweden AB.

## OpenCL

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

## Trademarks

NVIDIA, the NVIDIA logo, and cuBLAS, CUDA, CUDA Toolkit, cuDNN, DALI, DIGITS, DGX, DGX-1, DGX-2, DGX Station, DLProf, GPU, JetPack, Jetson, Kepler, Maxwell, NCCL, Nsight Compute, Nsight Systems, NVCAffe, NVIDIA Ampere GPU architecture, NVIDIA Deep Learning SDK, NVIDIA Developer Program, NVIDIA GPU Cloud, NVLink, NVSHMEM, PerfWorks, Pascal, SDK Manager, T4, Tegra, TensorRT, TensorRT Inference Server, Tesla, TF-TRT, Triton Inference Server, Turing, and Volta are trademarks and/or registered trademarks of NVIDIA Corporation in the United States and other countries. Other company and product names may be trademarks of the respective companies with which they are associated.

## Copyright

© 2016-2020 NVIDIA Corporation. All rights reserved.

