



# NVIDIA TensorRT

Support Matrix | NVIDIA Docs

# Table of Contents

Chapter 1. Features For Platforms And Software..... 1

Chapter 2. Layers And Features.....3

Chapter 3. Layers And Precision..... 7

Chapter 4. Hardware And Precision..... 10

Chapter 5. Layers For Flow-Control Constructs..... 12

Chapter 6. Compute Capability Per Platform..... 15

Chapter 7. Software Versions Per Platform..... 16

Chapter 8. ONNX Operator Support.....17

# Chapter 1. Features For Platforms And Software

This section lists the supported NVIDIA® TensorRT™ features based on which platform and software.

Table 1. List of supported features per platform.

	Linux x86-64	Windows x64	Linux ppc64le	Linux AArch64
	8.2.x	8.2.x	8.0.x	8.0.x
<a href="#">Supported NVIDIA CUDA® versions</a>	<a href="#">11.5</a> <sup>1</sup> <a href="#">11.4 update 3</a> <a href="#">11.3 update 1</a> <sup>1</sup> <a href="#">11.2 update 2</a> <sup>1</sup> <a href="#">11.1 update 1</a> <sup>1</sup> <a href="#">11.0 update 1</a> <sup>1</sup> <a href="#">10.2</a>	<a href="#">11.5</a> <sup>2</sup> <a href="#">11.4 update 3</a> <a href="#">11.3 update 1</a> <a href="#">11.2 update 2</a> <a href="#">11.1 update 1</a> <a href="#">11.0 update 1</a> <a href="#">10.2</a>	<a href="#">11.3 update 1</a>	<a href="#">11.4 update 3</a> <a href="#">10.2</a>
<a href="#">Supported cuBLAS versions</a>	11.7.3.x 11.6.5.x 11.5.1.109 11.4.1.1043 11.3.0.106 11.2.0.252 10.2.3.254	11.7.3.x 11.6.5.x 11.5.1.109 11.4.1.1043 11.3.0.106 11.2.0.252 10.2.3.254	11.5.1.109	11.6.5.x 10.2.2.214

<sup>1</sup> These CUDA versions are supported using a single build, built with CUDA toolkit 11.4. It is compatible with all CUDA 11.x versions and only requires driver 450.x.

<sup>2</sup> The NVRTC dynamic library from CUDA toolkit version 11.4 is required when using CUDA toolkits newer than 11.4.

	Linux x86-64	Windows x64	Linux ppc64le	Linux AArch64
	8.2.x	8.2.x	8.0.x	8.0.x
<a href="#">Supported cuDNN versions</a>	<a href="#">cuDNN 8.2.1</a>	<a href="#">cuDNN 8.2.1</a>	<a href="#">cuDNN 8.2.1</a>	<a href="#">cuDNN 8.2.1</a>
TensorRT Python API	Yes	Yes	Yes	Yes
NvUffParser	Yes	Yes	Yes	Yes
NvOnnxParser	Yes	Yes	Yes	Yes
<a href="#">Loops</a>	Yes	Yes	Yes	Yes

**Note:**

- Serialized engines are not portable across platforms or TensorRT versions.
- Refer to the minimum compatible driver versions in the [NVIDIA CUDA Release Notes](#) for specific [NVIDIA Driver](#) versions.

## Chapter 2. Layers And Features

The section lists the supported TensorRT layers and each of the features.

### About this task



#### Note:

- **Supports broadcast** indicates support for broadcast in this layer. This layer allows its two input tensors to be of dimensions [1, 5, 4, 3] and [1, 5, 1, 1], and its output is [1, 5, 4, 3]. The second input tensor has been broadcast in the innermost 2 dimensions.
- **Supports broadcast across batch** indicates support for broadcast across the batch dimension. "NA" in this column means it's not allowed in networks with an implicit batch dimension.

Table 2. List of supported features per TensorRT layer.

Layer	Dimensions of input tensor	Dimensions of output tensor	Does the operation apply to only the innermost 3 dimensions?	Supports broadcast	Supports broadcast across batch
<a href="#">IActivationLayer</a>	0-7 dimensions	0-7 dimensions	No	No	No
<a href="#">IAssertionLayer</a>	0-1 dimensions	No output	No	No	No
<a href="#">IConcatenationLayer</a>	1-7 dimensions	1-7 dimensions	No	No	No
<a href="#">IConditionLayer</a>	0	No output	No	No	No
<a href="#">IConstantLayer</a>	has no inputs	0-7 dimensions	No	No	Always
<a href="#">IConvolutionLayer</a> > 2D <a href="#">Convolution</a>	3 or more dimensions	3 or more dimensions	Yes	No	No

Layer	Dimensions of input tensor	Dimensions of output tensor	Does the operation apply to only the innermost 3 dimensions?	Supports broadcast	Supports broadcast across batch
<a href="#">IConvolutionLayer &gt; 3D Convolution</a>	4 or more dimensions	4 or more dimensions	No	No	No
<a href="#">IDeconvolutionLayer &gt; 2D Deconvolution</a>	3 or more dimensions	3 or more dimensions	Yes	No	No
<a href="#">IDeconvolutionLayer &gt; 3D Deconvolution</a>	4 or more dimensions	4 or more dimensions	No	No	No
<a href="#">IDequantizeLayer</a>	2 or more dimensions	2 or more dimensions	Yes	No	No
<a href="#">IEinsumLayer</a>	0-7 dimensions	0-7 dimensions	No	No	Yes
<a href="#">IElementWiseLayer</a>	0-7 dimensions	0-7 dimensions	No	Yes	Yes
<a href="#">IFillLayer</a>	1 dimension	0-7 dimensions	No	NA	NA
<a href="#">IFullyConnectedLayer</a>	3 or more dimensions	3 or more dimensions	Yes	No	No
<a href="#">IGatherLayer</a>	<ul style="list-style-type: none"> <li>▶ Input1: 1-7 dimensions</li> <li>▶ Input2: 0-7 dimensions</li> </ul>	0-7 dimensions	No	No	Yes
<a href="#">IIdentityLayer</a>	0-7 dimensions	0-7 dimensions	No	No	No
<a href="#">IIfConditionalOutputLayer</a>	0-7 dimensions	0-7 dimensions	No	No	No
<a href="#">IIfConditionalInputLayer</a>	0-7 dimensions	0-7 dimensions	No	No	No
<a href="#">IIteratorLayer</a>	1-7 dimensions	0-6 dimensions	No	No	NA
<a href="#">ILoopOutputLayer</a>	0-7 dimensions	0-7 dimensions	No	No	NA
<a href="#">ILRNLayer</a>	3 or more dimensions	3 or more dimensions	Yes	No	No
<a href="#">IMatrixMultiplyLayer</a>	2 or more dimensions	2 or more dimensions	No	Yes	Yes
<a href="#">IPaddingLayer</a>	3 or more dimensions	3 or more dimensions	Yes	No	No

Layer	Dimensions of input tensor	Dimensions of output tensor	Does the operation apply to only the innermost 3 dimensions?	Supports broadcast	Supports broadcast across batch
<a href="#">IParametricReLULayer</a>	1-7 dimensions	1-7 dimensions	No	No	No
<a href="#">IPluginV2Layer</a>	User defined	User defined	User defined	User defined	User defined
<a href="#">IPoolingLayer &gt; 2D Pooling</a>	3 or more dimensions	3 or more dimensions	Yes	Yes	Yes
<a href="#">IPoolingLayer &gt; 3D Pooling</a>	4 or more dimensions	4 or more dimensions	No	Yes	Yes
<a href="#">IQuantizeLayer</a>	2 or more dimensions	2 or more dimensions	Yes	No	No
<a href="#">IRaggedSoftMaxLayer</a>	<ul style="list-style-type: none"> <li>► Input: 2 dimensions</li> <li>► Bounds: 2 dimensions</li> </ul>	2 or more dimensions	No	No	Yes
<a href="#">IRecurrenceLayer</a>	0-7 dimensions	0-7 dimensions	No	No	NA
<a href="#">IReduceLayer</a>	1-7 dimensions	0-7 dimensions	No	No	No
<a href="#">IResizeLayer</a>	1-7 dimensions	1-7 dimensions	No	No	No
<a href="#">IRNNv2Layer</a>	<ul style="list-style-type: none"> <li>► Data/Hidden/Cell: 2 or more dimensions</li> <li>► SeqLen: 0 or more dimensions</li> </ul>	Data/Hidden/Cell: 2 or more dimensions	No	No	No
<a href="#">IScaleLayer</a>	3 or more dimensions	3 or more dimensions	Yes	No	No
<a href="#">IScatterLayer</a>	0-7 dimensions	0-7 dimensions	No	No	No
<a href="#">ISelectLayer</a>	0-7 dimensions	0-7 dimensions	No	Yes	NA
<a href="#">IShapeLayer</a>	1 or more dimensions	1 dimension	No	No	NA
<a href="#">IShuffleLayer</a>	0-7 dimensions	0-7 dimensions	No	No	No

Layer	Dimensions of input tensor	Dimensions of output tensor	Does the operation apply to only the innermost 3 dimensions?	Supports broadcast	Supports broadcast across batch
<a href="#">ISliceLayer</a>	1-7 dimensions	1-7 dimensions	No	No	Yes
<a href="#">ISoftMaxLayer</a>	1-7 dimensions	1-7 dimensions	No	No	Yes
<a href="#">ITopKLayer</a>	1-7 dimensions	<ul style="list-style-type: none"> <li>► Output1: 1-7 dimensions</li> <li>► Output2: 1-7 dimensions</li> </ul>	Yes	No	Yes
<a href="#">ITripLimitLayer</a>	0 dimensions	has no outputs	No	No	NA
<a href="#">IUnaryLayer</a>	1-7 dimensions	1-7 dimensions	No	No	No

For more information about each of the TensorRT layers, see [TensorRT Layers](#).



## Chapter 3. Layers And Precision

The section lists the TensorRT layers and the precision modes that each layer supports. It also lists the ability of the layer to run on Deep Learning Accelerator (DLA).

For more information about additional constraints, see [DLA Supported Layers](#).

For more information about each of the TensorRT layers, see [TensorRT Layers](#). To view a list of the specific attributes that are supported by each layer, refer to the [NVIDIA TensorRT API Reference](#) documentation.

Table 3. List of supported precision modes per TensorRT layer.

Layer	FP32	FP16	INT8	INT32	Bool	DLA FP16	DLA INT8
<a href="#">IActivationLayer</a>	Yes	Yes	Yes	No	No	Yes <sup>3</sup>	Yes <sup>4</sup>
<a href="#">IAssertionLayer</a>	No	No	No	No	Yes	No	No
<a href="#">IConcatenationLayer</a>	Yes	Yes	Yes	Yes	No	Yes <sup>5</sup>	Yes <sup>5</sup>
<a href="#">IConditionLayer</a>	No	No	No	No	Yes	No	No
<a href="#">IConstantLayer</a>	Yes	Yes	Yes	Yes	No	No	No
<a href="#">IConvolutionLayer &gt; 2D Convolution</a>	Yes	Yes	Yes	No	No	Yes	Yes
<a href="#">IConvolutionLayer &gt; 3D Convolution</a>	Yes	Yes	Yes	No	No	No	No
<a href="#">IDeconvolutionLayer &gt; 2D Deconvolution</a>	Yes	Yes	Yes	No	No	Yes	Yes <sup>6</sup>

<sup>3</sup> Partial support. Yes for ReLU, Clipped ReLU, Leaky ReLU, Sigmoid and TanH activation types only.

<sup>4</sup> Partial support. Yes for ReLU, Clipped ReLU, Leaky ReLU, Sigmoid and TanH activation types only.

<sup>5</sup> Partial support. Yes for concatenation across c dimension only.

<sup>6</sup> Partial support. Yes for ungrouped deconvolutions and No for grouped.

Layer	FP32	FP16	INT8	INT32	Bool	DLA FP16	DLA INT8
<a href="#">IDeconvolutionLayer &gt; 3D Deconvolution</a>	Yes	Yes	No	No	No	No	No
<a href="#">IDequantizeLayer</a>	No	No	Yes	No	No	No	No
<a href="#">IEinsumLayer</a>	Yes	Yes	No	No	No	No	No
<a href="#">IElementWiseLayer</a>	Yes	Yes	No	Yes	Yes	Yes <sup>7</sup>	Yes <sup>8</sup>
<a href="#">IFillLayer</a>	Yes	No	No	Yes	No	No	No
<a href="#">IFullyConnectedLayer</a>	Yes	Yes	Yes	No	No	Yes	Yes
<a href="#">IGatherLayer</a>	Yes	Yes	No	Yes	No	No	No
<a href="#">IIdentityLayer</a>	Yes	Yes	Yes	Yes	No	No	No
<a href="#">IfConditionalOutputLayer</a>	Yes	Yes	No	Yes	Yes	No	No
<a href="#">IfConditionalInputLayer</a>	Yes	Yes	No	Yes	Yes	No	No
<a href="#">IIteratorLayer</a>	Yes	Yes	No	Yes	No	No	No
<a href="#">ILoopOutputLayer</a>	Yes	Yes	No	Yes	No	No	No
<a href="#">ILRNLayer</a>	Yes	Yes	Yes	No	No	Yes	No
<a href="#">IMatrixMultiplyLayer</a>	Yes	Yes	No	No	No	No	No
<a href="#">IPaddingLayer</a>	Yes	Yes	Yes	No	No	No	No
<a href="#">IParametricReLULayer</a>	Yes	Yes	Yes	No	No	No	No
<a href="#">IPluginV2Layer</a>	Yes	Yes	Yes	No	No	No	No
<a href="#">IPoolingLayer &gt; 2D Pooling</a>	Yes	Yes	Yes	No	No	Yes <sup>9</sup>	Yes <sup>9</sup>
<a href="#">IPoolingLayer &gt; 3D Pooling</a>	Yes	Yes	No	No	No	No	No
<a href="#">IQuantizeLayer</a>	Yes	No	No	No	No	No	No
<a href="#">IRaggedSoftmaxLayer</a>	Yes	No	No	No	No	No	No
<a href="#">IRecurrenceLayer</a>	Yes	Yes	No	Yes	Yes	No	No
<a href="#">IReduceLayer</a>	Yes	Yes	Yes	Yes	No	No	No
<a href="#">IResizeLayer</a>	Yes	Yes	No	No	No	No	No
<a href="#">IRNNv2Layer</a>	Yes	Yes	No	No	No	No	No

<sup>7</sup> Partial support. Yes for sum, sub, prod, min and max elementwise operations only.

<sup>8</sup> Partial support. Yes for sum elementwise operation only.

<sup>9</sup> Partial support. Yes for max and average padding inclusive pooling type only.

Layer	FP32	FP16	INT8	INT32	Bool	DLA FP16	DLA INT8
<a href="#">IScaleLayer</a>	Yes	Yes	Yes	No	No	Yes <sup>10</sup>	Yes <sup>10</sup>
<a href="#">IScatterLayer</a>	Yes	Yes	Yes	Yes	No	No	No
<a href="#">ISelectLayer</a>	Yes	Yes	No	Yes	Yes	No	No
<a href="#">IShapeLayer</a>	Yes <sup>11</sup>	Yes	Yes	Yes	Yes	No	No
<a href="#">IShuffleLayer</a>	Yes	Yes	Yes	Yes	No	No	No
<a href="#">ISliceLayer</a>	Yes	Yes	No <sup>12</sup>	Yes	No	No	No
<a href="#">ISoftMaxLayer</a>	Yes	Yes	No	No	No	No	No
<a href="#">ITopKLayer</a>	Yes	Yes	No	No	No	No	No
<a href="#">ITripLimitLayer</a>	Yes	Yes	No	Yes	Yes	No	No
<a href="#">IUnaryLayer</a>	Yes	Yes	No	No	Yes	No	No



**Note:** DLA with FP16/INT8 precision with some restrictions on layer parameters.

<sup>10</sup> Partial support. DLA does not support power on scale layer.

<sup>11</sup> Output is always INT32.

<sup>12</sup> Partial support. *Yes* for unstrided Slice and *No* for strided.

## Chapter 4. Hardware And Precision

The following table lists NVIDIA hardware and which precision modes each hardware supports. TensorRT supports all NVIDIA hardware with capability SM 5.0 or higher. It also lists the availability of DLA on this hardware. Refer to the following tables for the specifics.



**Note:** Support for CUDA compute capability version 3.0 has been removed. Support for CUDA compute capability versions below 5.0 may be removed in a future release and is now deprecated.

Table 4. Supported hardware

<a href="#">CUDA Compute Capability</a>	Example Device	TF32	FP32	FP16	INT8	FP16 Tensor Cores	INT8 Tensor Cores	DLA
8.6	NVIDIA A10	Yes	Yes	Yes	Yes	Yes	Yes	No
8.0	NVIDIA A100/ GA100 GPU	Yes	Yes	Yes	Yes	Yes	Yes	No
7.5	Tesla T4	No	Yes	Yes	Yes	Yes	Yes	No
7.2	Jetson AGX Xavier	No	Yes	Yes	Yes	Yes	Yes	Yes
7.0	Tesla V100	No	Yes	Yes	Yes	Yes	No	No
6.2	Jetson TX2	No	Yes	Yes	No	No	No	No
6.1	Tesla P4	No	Yes	No	Yes	No	No	No
6.0	Tesla P100	No	Yes	Yes	No	No	No	No

<u>CUDA Compute Capability</u>	<u>Example Device</u>	TF32	FP32	FP16	INT8	FP16 Tensor Cores	INT8 Tensor Cores	DLA
5.3	Jetson TX1	No	Yes	Yes	No	No	No	No
5.2	Tesla M4	No	Yes	No	No	No	No	No
5.0	Quadro K2200	No	Yes	No	No	No	No	No

### Deprecated hardware

Table 5. List of supported precision mode per hardware.

<u>CUDA Compute Capability</u>	<u>Example Device</u>	FP32	FP16	INT8	FP16 Tensor Cores	INT8 Tensor Cores	DLA
3.7	Tesla K80	Yes	No	No	No	No	No
3.5	Tesla K40	Yes	No	No	No	No	No

### Removed hardware

Table 6. List of supported precision mode per hardware.

<u>CUDA Compute Capability</u>	<u>Example Device</u>	FP32	FP16	INT8	FP16 Tensor Cores	INT8 Tensor Cores	DLA
3.0	Tesla K10	Yes	No	No	No	No	No

---

## Chapter 5. Layers For Flow-Control Constructs

The following table lists the TensorRT layers that can be used as interior layers in TensorRT flow-control constructs.

Currently, TensorRT supports loop constructs (via `ILoopLayer`) and ternary conditional constructs (via `IIfConditionalLayer`). Interior layers are layers that comprise the body of a loop or one of the two branches of an if-conditional.

An `ILoopLayer` interior layer may contain other loops and/or if-conditionals. An `IIfConditionalLayer` branch may contain other if-conditionals and/or loops.

Flow-control constructs do not support INT8 calibration and interior-layers cannot employ implicit-quantization (INT8 is supported only in explicit-quantization mode).

Table 7. List of TensorRT layers that are supported as interior layers of flow-control constructs

Layer	Supported
<a href="#">IActivationLayer</a>	Yes, when the operation is one of: <code>kRELU</code> , <code>kSIGMOID</code> , <code>kTANH</code> , <code>kELU</code>
<a href="#">IAssertionLayer</a>	Yes
<a href="#">IConcatenationLayer</a>	Yes
<a href="#">IConditionLayer</a>	Yes (for nested conditionals)
<a href="#">IConstantLayer</a>	Yes
<a href="#">IConvolutionLayer &gt; 2D Convolution</a>	singleton channel and spatial dims, i.e. said dimensions must be static or have a single value in each optimization profile
<a href="#">IConvolutionLayer &gt; 3D Convolution</a>	singleton channel and spatial dims
<a href="#">IDeconvolutionLayer &gt; 2D Deconvolution</a>	No
<a href="#">IDeconvolutionLayer &gt; 3D Deconvolution</a>	No
<a href="#">IDequantizeLayer</a>	No

Layer	Supported
<a href="#">IEinsumLayer</a>	Yes
<a href="#">IElementWiseLayer</a>	Yes
<a href="#">IFillLayer</a>	kRANDOM_UNIFORM only
<a href="#">IFullyConnectedLayer</a>	Yes
<a href="#">IGatherLayer</a>	Yes
<a href="#">IIdentityLayer</a>	Yes
<a href="#">IIfConditionalOutputLayer</a>	Yes (for nested conditionals)
<a href="#">IIfConditionalInputLayer</a>	Yes (for nested conditionals)
<a href="#">IIteratorLayer</a>	Yes (for nested loops)
<a href="#">ILoopOutputLayer</a>	Yes (for nested loops)
<a href="#">ILRNLayer</a>	No
<a href="#">IMatrixMultiplyLayer</a>	Yes
<a href="#">IPaddingLayer</a>	No
<a href="#">IParametricReluLayer</a>	No
<a href="#">IPluginV2Layer</a>	Yes
<a href="#">IPoolingLayer &gt; 2D Pooling</a>	No
<a href="#">IPoolingLayer &gt; 3D Pooling</a>	No
<a href="#">IQuantizeLayer</a>	No
<a href="#">IRaggedSoftMaxLayer</a>	No
<a href="#">IRecurrenceLayer</a>	Yes
<a href="#">IReduceLayer</a>	Yes
<a href="#">IResizeLayer</a>	No
<a href="#">IRNNv2Layer</a>	No
<a href="#">IScaleLayer</a>	Yes
<a href="#">IScatterLayer</a>	Yes
<a href="#">ISelectLayer</a>	Yes
<a href="#">IShapeLayer</a>	Yes
<a href="#">IShuffleLayer</a>	Yes
<a href="#">ISliceLayer</a>	Yes
<a href="#">ISoftMaxLayer</a>	Yes
<a href="#">ITopKLayer</a>	No
<a href="#">ITripLimitLayer</a>	Yes

Layer	Supported
<a href="#">UnaryLayer</a>	Yes, when the operation is one of: <code>kABS</code> , <code>kCEIL</code> , <code>kERF</code> , <code>kEXP</code> , <code>kFLOOR</code> , <code>kLOG</code> , <code>kNEG</code> , <code>kNOT</code> , <code>kRECIP</code> , <code>kROUND</code> , <code>kSIGN</code> , <code>kSQRT</code> , <code>kSIN</code> , <code>kCOS</code> , <code>kATAN</code>



---

## Chapter 6. Compute Capability Per Platform

The section lists the supported compute capability based on platform.

Table 8. Compute capability per platform

Platform	Compute capability
Linux x86-64	3.5, 3.7, 5.0, 5.2, 6.0, 6.1, 7.0, 7.5, 8.0 <sup>13</sup> , 8.6 <sup>14</sup>
Windows 10 x64	3.5, 3.7, 5.0, 5.2, 6.0, 6.1, 7.0, 7.5, 8.0 <sup>13</sup> , 8.6 <sup>14</sup>
CentOS 8.3 ppc64le	7.0, 7.5, 8.0, 8.6
Ubuntu 20.04 SBSA	7.0, 7.5, 8.0, 8.6
JetPack AArch64	5.3, 6.2, 7.2

---

<sup>13</sup> Requires CUDA toolkit 11.0 or newer and a TensorRT CUDA 11.x build.

<sup>14</sup> Requires CUDA toolkit 11.1 or newer and a TensorRT CUDA 11.x build.

---

## Chapter 7. Software Versions Per Platform

The section lists the supported software versions based on platform.

Table 9. List of supported platforms per software version.

Platform	Compiler version	Python version
Ubuntu 18.04 x86-64	<a href="#">gcc 8.3.1</a>	<a href="#">3.6</a>
Ubuntu 20.04 x86-64	<a href="#">gcc 8.3.1</a>	<a href="#">3.8</a>
CentOS 7.9 x86-64	<a href="#">gcc 8.3.1</a>	<a href="#">3.6</a>
CentOS 8.3 x86-64	<a href="#">gcc 8.3.1</a>	<a href="#">3.8</a>
SLES 15 x86-64	<a href="#">gcc 8.3.1</a>	N/A
Windows 10 x64	<a href="#">MSVC 2017u5</a>	N/A
CentOS 8.3 ppc64le	<a href="#">Clang 10.0.1</a>	<a href="#">3.8</a>
Ubuntu 20.04 SBSA	<a href="#">gcc 8.4.0</a>	<a href="#">3.8</a>
JetPack AArch64	<a href="#">gcc 7.5.0</a>	<a href="#">3.6</a>



**Note:** Python versions supported when using Debian or RPM packages. When using Python wheel files, versions 3.6, 3.7, 3.8, and 3.9 are supported.

---

## Chapter 8. ONNX Operator Support

The ONNX operator support list for TensorRT can be found [here](#).

## 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.

## ARM

ARM, AMBA and ARM Powered are registered trademarks of ARM Limited. Cortex, MPCore and Mali are trademarks of ARM Limited. "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.

## HDMI

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

## BlackBerry/QNX

Copyright © 2020 BlackBerry Limited. All rights reserved.

Trademarks, including but not limited to BLACKBERRY, EMBLEM Design, QNX, AVIAGE, MOMENTICS, NEUTRINO and QNX CAR are the trademarks or registered trademarks of BlackBerry Limited, used under license, and the exclusive rights to such trademarks are expressly reserved.

## Google

Android, Android TV, Google Play and the Google Play logo are trademarks of Google, Inc.

## Trademarks

NVIDIA, the NVIDIA logo, and CUDA, DALI, DRIVE, JetPack, Jetson AGX Xavier, Jetson Nano, Kepler, Maxwell, NGC, Nsight, Pascal, Quadro, Tegra, TensorRT, Triton, 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

© 2018-2021 NVIDIA Corporation & affiliates. All rights reserved.