



# NVIDIA DATA LOADING LIBRARY (DALI)

RN-09096-001 \_v01 | December 2019

## Release Notes



# TABLE OF CONTENTS

Chapter 1. DALI Overview.....	1
Chapter 2. DALI Release 0.16.0 Beta.....	2
Chapter 3. DALI Release 0.15.0 Beta.....	5
Chapter 4. DALI Release 0.14.0 Beta.....	7
Chapter 5. DALI Release 0.13.0 Beta.....	9
Chapter 6. DALI Release 0.12.0 Beta.....	11
Chapter 7. DALI Release 0.11.0 Beta.....	13
Chapter 8. DALI Release 0.10.0 Beta.....	15
Chapter 9. DALI Release 0.9.1 Beta.....	17
Chapter 10. DALI Release 0.8.1 Beta.....	19
Chapter 11. DALI Release 0.8 Beta.....	21
Chapter 12. DALI Release 0.7 Beta.....	23
Chapter 13. DALI Release 0.6.1 Beta.....	25
Chapter 14. DALI Release 0.6 Beta.....	27
Chapter 15. DALI Release 0.5 Beta.....	29
Chapter 16. DALI Release 0.4.1 Beta.....	30
Chapter 17. DALI Release 0.4 Beta.....	31
Chapter 18. DALI Release 0.3 Beta.....	33
Chapter 19. DALI Release 0.2 Beta.....	35
Chapter 20. DALI Release 0.1.2 Beta.....	37
Chapter 21. DALI Release 0.1.1 Beta.....	38

# Chapter 1.

## DALI OVERVIEW

NVIDIA® Data Loading Library™ (DALI) accelerates the preprocessing of input data for deep learning applications. By offloading augmentations onto GPUs, DALI addresses performance bottlenecks in today's computer vision deep learning applications that include complex, multi-stage data augmentation steps. With DALI beta release, deep learning researchers can scale training performance on image classification models such as ResNet-50 with MXNet, TensorFlow, and PyTorch across Amazon Web Services P3 8 GPU instances or DGX-1 systems with Volta architecture. Framework developers will have less duplication due to better code reuse and maintainability.

DALI offers both performance and flexibility of accelerating different data pipelines (graphs that can have multiple outputs and inputs), as a single library, that can be easily integrated into different deep learning training and inference applications.

# Chapter 2.

## DALI RELEASE 0.16.0 BETA

The DALI 0.16.0 is a beta release, therefore, all features, functionality, and performance will likely be limited.

### Using DALI 0.16.0 Beta

To upgrade to DALI 0.16.0 beta from an older version of DALI, follow the installation instructions in the [DALI Quick Start Guide](#).

Refer to the [DALI Developer Guide](#) for usage details.



The internal DALI C++ API used for operators implementation, and the C++ API that enables using DALI as a library from native code, are not yet officially supported. Hence these APIs may change in the next release without advance notice.

### Key Features and Enhancements

This DALI release includes the following key features and enhancements.

- ▶ SupportOps are removed and CPU operators can now be used in their place ([#1423](#))
  - ▶ Allows for more generic shapes to be provided as argument inputs
  - ▶ Allows to return the data previously generated by support Op (like `ops.Uniform`) from pipeline
  - ▶ Internally Argument Inputs accept TensorVector instead of Tensor.
- ▶ Extended support for TensorFlow dataset API:
  - ▶ GPU ([#1354](#)) and initial multi-GPU ([#1382](#)) support
  - ▶ TensorFlow 1.15 and 2.0 support for TensorFlow dataset ([#1395](#))
- ▶ New operators:
  - ▶ HSV manipulation operator for GPU and CPU ([#1338](#))
  - ▶ GPU `dltensor` operator ([#1261](#))

- ▶ Pad operator ([#1180](#))
- ▶ Initial support for arithmetic operators ([#1322](#), [#1355](#), [#1426](#), [#1449](#)):
  - ▶ Regular Python arithmetic expressions inside DALI pipeline definition
  - ▶ Supported binary operators: +, -, \*, /, //
- ▶ Improvements to current operators:
  - ▶ Support float16 in Cast GPU operator ([#1368](#))
  - ▶ Rotate operator using a new Warp kernel ([#1403](#))
  - ▶ Optimized implementation for **WarpAffine** ([#1387](#))
  - ▶ Added **additional\_decode\_surfaces** parameter to videoreader ([#1393](#))
- ▶ New and updated examples:
  - ▶ **BrightnessContrast** operator example ([#1414](#))
  - ▶ New warp example ([#1158](#))
- ▶ Build system improvements for architectures other than x86
- ▶ Operators moved to separate library ([#1380](#), [#1384](#), [#1406](#))

## Fixed Issues

This DALI release includes the following fixes.

- ▶ Fixed DALI TensorFlow plugin CXX11 ABI issue ([#1361](#))
- ▶ Fixed DALI TensorFlow installation for TensorFlow 2.0 ([#1386](#))

## Breaking Changes

- ▶ Support operators were removed and their usage replaced with CPU operators. Old support operators were converted to CPU operators.
- ▶ DALI operators were moved into a separate library to detach them from the pipeline. It is linked into the target DALI library.

## Deprecated Features

- ▶ Removed the following deprecated operators:
  - ▶ **HostDecoder** and **nvJPEGDecoder** (generic **ImageDecoder** should be used instead) ([#1398](#))
  - ▶ **NormalizePermute** (**CropMirrorNormalize** should be used instead) ([#1402](#))
  - ▶ DALI 0.17 is the last official release for Python 2.7, which reaches the end of life on January 1st, 2020. To stay up to date with DALI, upgrade to Python 3.5 or later.
  - ▶ **Crop**, **CropMirrorNormalize** and **Slice** operator possible output types are limited to one of **uint8\_t**, **int16\_t**, **uint16\_t**, **int32\_t**, **float**, **float16** or passing through the input type ([#1418](#)).

## Known Issues

- ▶ The video loader operator requires that the key frames occur at a minimum every 10 to 15 frames of the video stream. If the key frames occur at a lesser frequency, then the returned frames may be out of sync.
- ▶ DALI TensorFlow plugin may not be compatible with TensorFlow versions 1.15.0 and/or later. If the user wants to use DALI with TensorFlow version which doesn't have prebuilt plugin binary shipped with DALI it requires the gcc compiler that matches the one used to build TensorFlow (gcc 4.8.4 or gcc, 4.8.5 or 5.4, depending on the particular version) is present on the system.
- ▶ Due to some known issues with meltdown/spectra mitigations and DALI, DALI shows best performance when run in Docker with escalated privileges, for example:
  - ▶ **privileged=yes** in Extra Settings for AWS data points
  - ▶ **--privileged** or **--security-opt seccomp=unconfined** for bare Docker

# Chapter 3.

## DALI RELEASE 0.15.0 BETA

The DALI 0.15.0 is a beta release. Hence, for all the features, the functionality and performance will likely be limited.

### Using DALI 0.15.0 Beta

To upgrade to DALI 0.15.0 beta from an older version of DALI, follow the installation instructions in the [DALI Quick Start Guide](#).

Refer to the [DALI Developer Guide](#) for usage details.



The internal DALI C++ API used for operators implementation, and the C++ API that enables using DALI as a library from native code, are not yet officially supported. Hence these APIs may change in the next release without advance notice.

### Key Features and Enhancements

This DALI release includes the following key features and enhancements.

- ▶ New operators and extended functionality of the following:
  - ▶ CPU **DLTensor** operator (#1233)
  - ▶ Added Lookup table operator (#1251)
  - ▶ **BrightnessContrast** operator (#1188)
  - ▶ Added **CropMirrorNormalize** 3D support (#1326)
  - ▶ Added **Reshape** operator (#1327)
- ▶ New kernels:
  - ▶ Linear transformation GPU kernel (#1262)
  - ▶ Linear Transformation kernel for CPU (#1300)
- ▶ Video reader improvements:
  - ▶ Improved heuristic for variable frame rate detection (#1242)

- ▶ Added support for decoding multiple resolution videos in the same pipeline ([#1144](#))
- ▶ Tensor layout ([#1237](#)) and Layout refactor ([#1250](#)) - new, domain agnostic, flexible description of a tensor layout that is compatible with string representation.
- ▶ Multi channel image decoding - supporting more than 3 channels:
  - ▶ `ImageDecoder libtiff` implementation ([#1264](#))
  - ▶ Handle `nchannels>3` in `ImageDecoder` ([#1285](#))
  - ▶ Parse number of channels in `PNGImage::PeekShape` ([#1288](#))
  - ▶ `TiffImage::PeekShapeImpl` parse and return number of channels ([#1304](#))
  - ▶ Add parsing of number of channels in `JpegImage::PeekShapeImpl` ([#1306](#))
- ▶ DALI TensorFlow plugin installation improvements.

### Fixed Issues

This DALI release includes the following fixes.

- ▶ Fixed `Transpose` operator when data shape with dimension of size 1 ([#1244](#))
- ▶ Fixed problems with seeking when stream `start_time` is `!= 0`. ([#1287](#))
- ▶ Suppressed warning when `FileReader` encounters dot and dot-dot entries ([#1318](#))
- ▶ Fixed `last_batch_padded` docs ([#1314](#))
- ▶ Fixed DALI TensorFlow plugin CXX11 ABI issue ([#1361](#))

### Breaking Changes

- ▶ Changed `Outputs` and `SharedOutputs` return type to tuple ([#1243](#))

### Known Issues

- ▶ The new video reader operator requires NVIDIA VIDEO CODEC SDK support in the platform. Prior to 19.01, the NVIDIA GPU Cloud (NGC) optimized containers lack this functionality in the default configuration. To enable the functionality, run the container with the `video` capability enabled, as shown below:

```
-e "NVIDIA_DRIVER_CAPABILITIES=compute,utility,video"
```

- ▶ The video loader operator requires that the key frames occur at a minimum every 10 to 15 frames of the video stream. If the key frames occur at a lesser frequency, then the returned frames may be out of sync.
- ▶ DALI TensorFlow plugin may not be compatible with TensorFlow versions 1.15.0 and/or later. If the user wants to use DALI with TensorFlow version which doesn't have prebuilt plugin binary shipped with DALI it requires the gcc compiler that matches the one used to build TensorFlow (gcc 4.8.4 or gcc, 4.8.5 or 5.4, depending on the particular version) is present on the system.



# Chapter 4.

## DALI RELEASE 0.14.0 BETA

### Using DALI 0.14.0 Beta

The DALI 0.14.0 can be used with the 19.09 NVIDIA GPU Cloud (NGC) optimized container for MXNet, PyTorch, and TensorFlow.

To upgrade to DALI 0.14.0 beta from an older version of DALI, follow the installation instructions in the [DALI Quick Start Guide](#).

Refer to the [DALI Developer Guide](#) for usage details.



The internal DALI C++ API used for operators implementation, and the C++ API that enables using DALI as a library from native code, are not yet officially supported. Hence these APIs may change in the next release without advance notice.

### Key Features and Enhancements

This DALI release includes the following key features and enhancements.

- ▶ Added an example for using DALI with multiple GPUs.
- ▶ Added a **Shapes** operator which returns the input shape as a tensor.
- ▶ Added conda recipes for DALI used for pre-built packages in Watson Machine Learning Community Edition (IBM Power and x86 systems).
- ▶ Extended the **ExternalSource** operator so it can hold more than one batch of inputs at a time.

### Fixed Issues

This DALI release includes the following fixes.

- ▶ Fixed FP16 bug from [#1129](#) and added FP16 test case. ([#1160](#))
- ▶ Fixed framework iterators behavior when **iter\_setup** raises **StopIteration**. ([#1136](#))

- ▶ Fixed `nvjpeg` legacy API. ([#1179](#))
- ▶ Fixed conversions to INT64 and UINT64. ([#1205](#))
- ▶ Fixed DALI TensorFlow install for conda environments. ([#1214](#))

### Breaking Changes

- ▶ Extended external source operator capacity ([#1127](#)) - it now requires input to be set for every iteration.
- ▶ Adjusted `Operator::Run` to take reference instead of pointer ([#1168](#)) (C++ Backend API).

### Deprecated Features

- ▶ Python `.cpu()` function for `EdgeReference` was removed to reduce confusion. ([#1181](#))

### Known Issues

- ▶ The new video reader operator requires NVIDIA VIDEO CODEC SDK support in the platform. Prior to 19.01, the NVIDIA GPU Cloud (NGC) optimized containers lack this functionality in the default configuration. To enable the functionality, run the container with the "video" capability enabled, as shown below:

```
-e "NVIDIA_DRIVER_CAPABILITIES=compute,utility,video"
```

- ▶ The video loader operator requires that the key frames occur at a minimum every 10 to 15 frames of the video stream. If the key frames occur at a lesser frequency, then the returned frames may be out of sync.
- ▶ DALI TensorFlow plugin may be not compatible with TensorFlow 1.14.0 release. The DALI TensorFlow plugin requires that the `gcc` compiler that matches the one used to build TensorFlow (`gcc 4.8.4` or `gcc 4.8.5`, depending on the particular version) be present on the system.

# Chapter 5.

## DALI RELEASE 0.13.0 BETA

The DALI 0.13.0 is a beta release. Hence, for all the features, the functionality and performance will likely be limited.

### Using DALI 0.13.0 Beta

The DALI 0.13.0 can be used with the 19.08 NVIDIA GPU Cloud (NGC) optimized container for MXNet, PyTorch, and TensorFlow. Also, the 19.08 container will be shipped with DALI 0.13.0.

To upgrade to DALI 0.13.0 beta from an older version of DALI, follow the installation instructions in the [DALI Quick Start Guide](#).

Refer to the [DALI Developer Guide](#) for usage details.



The internal DALI C++ API used for operators implementation, and the C++ API that enables using DALI as a library from native code, are not yet officially supported. Hence these APIs may change in the next release without advance notice.

### Key Features and Enhancements

This DALI release includes the following key features and enhancements.

- ▶ Added new:
  - ▶ Fast coco reader
  - ▶ TorchPythonFunction operator
  - ▶ Sink operators
- ▶ Reworked how the reader pick samples from the shuffling buffer.
- ▶ Used DALI\_extra for test data.
- ▶ Added checks to see if Python API is not mixed between simple, scheduled and iterator APIs.
- ▶ Added support for reading video files with labels using `file_list` argument.

## Fixed Issues

This DALI release includes the following fixes.

- ▶ Restored support of `use_batched_decode` argument in `nvJPEGDecoder` operator (only for legacy `nvJPEGDecoder` implementation).
- ▶ Fixed FP16 support in DALI TensorFlow plugin.
- ▶ Fixed Python operator with side effects.
- ▶ Fixed a race condition in async pipeline executor.
- ▶ Disabled `video_reader_op` test when NVDEC is disabled.
- ▶ Fixed sampling of chroma in the VideoReader operation.
- ▶ Fix detection pipeline example.

## Breaking Changes

- ▶ Reader sampling from shuffling buffer was adjusted. Now samples are not mixed between epochs.

## Deprecated Features

- ▶ Deprecated `NormalizePermute` in favor of `CropMirrorNormalize`
- ▶ Multiple Input Sets handling was removed from backend and is only Python level syntactic sugar.

## Known Issues

- ▶ The new video reader operator requires NVIDIA VIDEO CODEC SDK support in the platform. Prior to 19.01, the NVIDIA GPU Cloud (NGC) optimized containers lack this functionality in the default configuration. To enable the functionality, run the container with the "video" capability enabled, as shown below:

```
-e "NVIDIA_DRIVER_CAPABILITIES=compute,utility,video"
```

- ▶ The video loader operator requires that the key frames occur at a minimum every 10 to 15 frames of the video stream. If the key frames occur at a lesser frequency, then the returned frames may be out of sync.
- ▶ DALI TensorFlow plugin may be not compatible with TensorFlow 1.14.0 release. The DALI TensorFlow plugin requires that the `gcc` compiler that matches the one used to build TensorFlow (`gcc 4.8.4` or `gcc 4.8.5`, depending on the particular version) be present on the system.

# Chapter 6.

## DALI RELEASE 0.12.0 BETA

The DALI 0.12.0 is a beta release. Hence, for all the features, the functionality and performance will likely be limited.

### Using DALI 0.12.0 Beta

The DALI 0.12.0 can be used with the 19.08 NVIDIA GPU Cloud (NGC) optimized container for MXNet, PyTorch, and TensorFlow. Also, the 19.08 container will be shipped with DALI 0.12.0.

To upgrade to DALI 0.12.0 beta from an older version of DALI, follow the installation instructions in the [DALI Quick Start Guide](#).

Refer to the [DALI Developer Guide](#) for usage details.



The internal DALI C++ API used for operators implementation, and the C++ API that enables using DALI as a library from native code, are not yet officially supported. Hence these APIs may change in the next release without advance notice.

### Key Features and Enhancements

This DALI release includes the following key features and enhancements.

- ▶ Added a new optimized implementation for [CropMirrorNormalize](#) for CPU and GPU.
- ▶ Added the ability to read video files together with labels.

### Fixed Issues

This DALI release includes the following fixes.

- ▶ Resumed shipping the prebuilt version of the TensorFlow plugin as a part of `nvidia-dali-tf-plugin` (as a result of fixing a compatibility issue). With this fix, the user must first install `nvidia-dali-tf-plugin` package. When this package

is installed, one of the prebuilt binaries of the TensorFlow plugin is selected. See [Binary Installation](#).

## Breaking Changes

- ▶ Source code build now requires the compiler be at least C++14 capable.

## Deprecated Features

- ▶ Deprecated the following as noted below:
  - ▶ `_run`: use `schedule_run` instead
  - ▶ `_share_outputs`: use `share_outputs` instead, and
  - ▶ `_release_outputs`: use `release_outputs` instead.
- ▶ Replaced `HostDecoder` and `nvJPEGDecoder` with generic [ImageDecoder](#). `ImageDecoder` is the recommended function for the image decoding, and the replaced APIs `HostDecoder` and `nvJPEGDecoder` will be removed in the future.
- ▶ `NormalizePermute` is replaced by the new [CropMirrorNormalize](#).

## Known Issues

- ▶ The new video reader operator requires NVIDIA VIDEO CODEC SDK support in the platform. Prior to 19.01, the NVIDIA GPU Cloud (NGC) optimized containers lack this functionality in the default configuration. To enable the functionality, run the container with the "video" capability enabled, as shown below:

```
-e "NVIDIA_DRIVER_CAPABILITIES=compute,utility,video"
```

- ▶ The video loader operator requires that the key frames occur at a minimum every 10 to 15 frames of the video stream. If the key frames occur at a lesser frequency, then the returned frames may be out of sync.
- ▶ DALI TensorFlow plugin may be not compatible with TensorFlow 1.14.0 release. The DALI TensorFlow plugin requires that the `gcc` compiler that matches the one used to build TensorFlow (`gcc 4.8.4` or `gcc 4.8.5`, depending on the particular version) be present on the system.

# Chapter 7.

## DALI RELEASE 0.11.0 BETA

The DALI 0.11.0 is a beta release. Hence, for all the features, the functionality and performance will likely be limited.

### Using DALI 0.11.0 Beta

The DALI 0.11.0 can be used with the 19.07 NVIDIA GPU Cloud (NGC) optimized container for MXNet, PyTorch, and TensorFlow. Also, the 19.07 container will be shipped with DALI 0.11.0.

To upgrade to DALI 0.11.0 beta from an older version of DALI, follow the installation instructions in the [DALI Quick Start Guide](#).

Refer to the [DALI Developer Guide](#) for usage details.



The internal DALI C++ API used for operators implementation, and the C++ API that enables using DALI as a library from native code, are not yet officially supported. Hence these APIs may change in the next release without advance notice.

### Key Features and Enhancements

This DALI release includes the following key features and enhancements.

- ▶ Added the ability to provide more than one input to, and return more than one output from, a Python-based operator.
- ▶ Extended the bounding box encoder for SSD to return offsets also. See [nvidia.dali.ops.BoxEncoder](#).
- ▶ Added the ability to build DALI by mounting the source code to the Docker so consecutive rebuilds are much faster.
- ▶ Added experimental support for **aarch64** (ARM) platform. Note that this support is added only for the native part—Python is not supported yet.
- ▶ Re-implemented the flip operator to increase its performance. See [nvidia.dali.ops.Flip](#).

- ▶ Improved the performance of nvJPEG Decoder with new internal API to match the previous implementation. See [nvidia.dali.ops.nvJPEGDecoder](https://nvidia.github.io/dali/ops/nvJPEGDecoder).

## Fixed Issues

This DALI release includes the following fixes.

- ▶ Fixed an issue wherein loading a plugin could break other operators that were already registered.
- ▶ Fixed an interoperability issue with PyCUDA. Now DALI will not interfere with PyCUDA's CUDA context management.

## Breaking API Changes

- ▶ CPU operators have moved from per-sample processing (pipeline process sample after sample, all the way through the pipeline) to batch-processing (all samples are processed by the first operator before moving to the next operator). This may result in a small performance degradation for some use cases. However, in the long term it will make some currently unavailable optimizations possible, for example: operations that need to view the whole batch during the processing (like random sample blending inside a batch).
- ▶ **CropCastPermute** is removed. **CropMirrorNormalize** should be used instead (with the default values for normalization).

## Deprecated Features

- ▶ Removed the prebuilt version of TensorFlow plugin for DALI. Now it is always necessary to install a separate **nvidia-dali-tf-plugin** package. See [Binary Installation](#).

## Known Issues

- ▶ The new video reader operator requires NVIDIA VIDEO CODEC SDK support in the platform. Prior to 19.01, the NVIDIA GPU Cloud (NGC) optimized containers lack this functionality in the default configuration. To enable the functionality, run the container with the "video" capability enabled, as shown below:
 

```
-e "NVIDIA_DRIVER_CAPABILITIES=compute,utility,video"
```
- ▶ The video loader operator requires that the key frames occur at a minimum every 10 to 15 frames of the video stream. If the key frames occur at a lesser frequency, then the returned frames may be out of sync.
- ▶ DALI TensorFlow plugin may be not compatible with TensorFlow 1.14.0 release. The DALI TensorFlow plugin requires that the **gcc** compiler that matches the one used to build TensorFlow (**gcc** 4.8.4 or **gcc** 4.8.5, depending on the particular version) be present on the system.



# Chapter 8.

## DALI RELEASE 0.10.0 BETA

The DALI 0.10.0 is a beta release. Hence, for all the features, the functionality and performance will likely be limited.

### Using DALI 0.10.0 Beta

The DALI 0.10.0 can be used with the 19.06 NVIDIA GPU Cloud (NGC) optimized container for MXNet, PyTorch, and TensorFlow. Also, the 19.06 container will be shipped with DALI 0.10.0.

To upgrade to DALI 0.10.0 beta from an older version of DALI, follow the installation instructions in the [DALI Quick Start Guide](#).

Refer to the [DALI Developer Guide](#) for usage details.



The internal DALI C++ API used for operators implementation, and the C++ API that enables using DALI as a library from native code, are not yet officially supported. Hence these APIs may change in the next release without advance notice.

### Key Features and Enhancements

This DALI release includes the following key features and enhancements.

- ▶ Reduced peak memory consumption. DALI tends to do new allocation before releasing the old memory during buffer resize. As it does not copy the old memory content, the old memory can be freed before allocating the new memory.
- ▶ Started publishing DALI nightly builds for CUDA 9 and CUDA 10, and weekly for CUDA 10.
- ▶ Added Python function operator. Now the user can create a Python-based operator that accepts one input and produces one output.

## Breaking API Changes

- ▶ None.

## Deprecated Features

- ▶ None.

## Known Issues

- ▶ The new video reader operator requires NVIDIA VIDEO CODEC SDK support in the platform. Prior to 19.01, the NVIDIA GPU Cloud (NGC) optimized containers lack this functionality in the default configuration. To enable the functionality, run the container with the "video" capability enabled, as shown below:

```
-e "NVIDIA_DRIVER_CAPABILITIES=compute,utility,video"
```

- ▶ The video loader operator requires that the key frames occur at a minimum every 10 to 15 frames of the video stream. If the key frames occur at a lesser frequency, then the returned frames may be out of sync.

# Chapter 9.

## DALI RELEASE 0.9.1 BETA

The DALI 0.9.1 is a beta release. Hence, for all the features, the functionality and performance will likely be limited.

### Using DALI 0.9.1 Beta

The DALI 0.9.1 can be used with the 19.05 NVIDIA GPU Cloud (NGC) optimized container for MXNet, PyTorch, and TensorFlow. Also, the 19.05 container will be shipped with DALI 0.9.1.

To upgrade to DALI 0.9.1 beta from an older version of DALI, follow the installation instructions in the [DALI Quick Start Guide](#).

Refer to the [DALI Developer Guide](#) for usage details.



The internal DALI C++ API used for operators implementation, and the C++ API that enables using DALI as a library from native code, are not yet officially supported. Hence these APIs may change in the next release without advance notice.

### Key Features and Enhancements

This DALI release includes the following key features and enhancements.

- ▶ Added Optical Flow example.
- ▶ Added full support for .pnm (.ppm/.pgm/.pnm) (in case of nvJPEG, fallback to CPU is done).
- ▶ Added the ability to do the lazy initialization of DALI. It is now possible to separate touching the data during the actual run, and at the time of all the necessary allocations during the build of the pipeline.
- ▶ Reduced the resize operator scratch buffer size by processing the data in mini batches.
- ▶ Added the ability to set the DALI CUDA streams priority.
- ▶ Fixed the sync issue in the DALI Python iterators.

- ▶ Fixed the initialization of CUDA context on the default device during pipeline creation.

### Breaking API Changes

- ▶ Internal Python pipeline API has changed. If any function `_*` was used, then that function should be updated to reflect the new semantic.

### Deprecated Features

- ▶ None.

### Known Issues

- ▶ The new video reader operator requires NVIDIA VIDEO CODEC SDK support in the platform. Prior to 19.01, the NVIDIA GPU Cloud (NGC) optimized containers lack this functionality in the default configuration. To enable the functionality, run the container with the "video" capability enabled, as shown below:

```
-e "NVIDIA_DRIVER_CAPABILITIES=compute,utility,video"
```

- ▶ The video loader operator requires that the key frames occur at a minimum every 10 to 15 frames of the video stream. If the key frames occur at a lesser frequency, then the returned frames may be out of sync.

# Chapter 10.

## DALI RELEASE 0.8.1 BETA

The DALI 0.8.1 is a beta release. Hence, for all the features, the functionality and performance will likely be limited.

### Using DALI 0.8.1 Beta

The DALI 0.8.1 can be used with the 19.04 NVIDIA GPU Cloud (NGC) optimized container for MXNet, PyTorch, and TensorFlow. Also, the 19.05 container will be shipped with DALI 0.9.

To upgrade to DALI 0.8.1 beta from an older version of DALI, follow the installation instructions in the [DALI Quick Start Guide](#).

Refer to the [DALI Developer Guide](#) for usage details.



The internal DALI C++ API used for operators implementation, and the C++ API that enables using DALI as a library from native code, are not yet officially supported. Hence these APIs may change in the next release without advance notice.

### Key Features and Enhancements

This DALI release includes the following key features and enhancements.

- ▶ Fixed nvJPEGDecoder cache when using the new nvJPEG decoupled API.

### Breaking API Changes

- ▶ None.

### Deprecated Features

- ▶ None.

## Known Issues

- ▶ The new video reader operator requires NVIDIA VIDEO CODEC SDK support in the platform. Prior to 19.01, the NVIDIA GPU Cloud (NGC) optimized containers lack this functionality in the default configuration. To enable the functionality, run the container with the "video" capability enabled, as below:

```
-e "NVIDIA_DRIVER_CAPABILITIES=compute,utility,video"
```

- ▶ The video loader operator requires that the key frames occur at a minimum every 10 to 15 frames of the video stream. If the key frames occur at a lesser frequency, then the returned frames may be out of sync.

# Chapter 11.

## DALI RELEASE 0.8 BETA

The DALI 0.8 is a beta release. Hence, for all the features, the functionality and performance will likely be limited.

### Using DALI 0.8 Beta

The DALI 0.8 can be used with the 19.03 NVIDIA GPU Cloud (NGC) optimized container for MXNet, PyTorch, and TensorFlow. Also, the 19.04 container will be shipped with DALI 0.8.

To upgrade to DALI 0.8 beta from an older version of DALI, follow the installation instructions in the [DALI Quick Start Guide](#).

Refer to the [DALI Developer Guide](#) for usage details.



The internal DALI C++ API used for operators implementation, and the C++ API that enables using DALI as a library from native code, are not yet officially supported. Hence these APIs may change in the next release without advance notice.

### Key Features and Enhancements

This DALI release includes the following key features and enhancements.

- ▶ Provides the CUDA 10-based DALI prebuilt binary.
- ▶ Added the customizable and decoupled executor stages. The CPU and GPU prefetch queues can now have different depths.
- ▶ The nvJPEG decoder now supports ROI-based decoding.
- ▶ The nvJPEG decoder now has split variants. The CPU and GPU part can be concurrently executed and independently prefetched.
- ▶ Improved the resize operator to support bilinear and triangular resampling. This improves the quality during aggressive down-scaling.
- ▶ Added the ability to return sparse tensor on CPU for TF DALI op.

- ▶ Fixed bugs in the detection pipeline operators. These bugs might have lead to inferior training accuracy.

### Breaking API Changes

- ▶ None.

### Known Issues

- ▶ The new video reader operator requires NVIDIA VIDEO CODEC SDK support in the platform. Prior to 19.01, the NVIDIA GPU Cloud (NGC) optimized containers lack this functionality in the default configuration. To enable the functionality, run the container with the "video" capability enabled, as below:

```
-e "NVIDIA_DRIVER_CAPABILITIES=compute,utility,video"
```

- ▶ The video loader operator requires that the key frames occur at a minimum every 10 to 15 frames of the video stream. If the key frames occur at a lesser frequency, then the returned frames may be out of sync.



# Chapter 12.

## DALI RELEASE 0.7 BETA

The DALI 0.7 is a beta release. Hence, for all the features, the functionality and performance will likely be limited.

### Using DALI 0.7 Beta

The DALI 0.7 can be used with the 19.03 NVIDIA GPU Cloud (NGC) optimized container for MXNet, PyTorch, and TensorFlow.

To upgrade to DALI 0.7 beta from an older version of DALI, follow the installation instructions in the [DALI Quick Start Guide](#).

Refer to the [DALI Developer Guide](#) for usage details.



The internal DALI C++ API used for operators implementation, and the C++ API that enables using DALI as a library from native code, are not yet officially supported. Hence these APIs may change in the next release without advance notice.

### Key Features and Enhancements

This DALI release includes the following key features and enhancements.

- ▶ Added new operators:
  - ▶ HostDecoderRandomCrop
  - ▶ Element Extract for sequences
- ▶ Added the API stub for incoming Optical Flow video operator.
- ▶ Introduced Python 3.7 based **whl** build.
- ▶ Added Docker-based "one click" build script.
- ▶ Added the ability to fine-tune the internal DALI buffers using per-operator presize hints to stage the output queues.
- ▶ Added **step**, **stride** and **shuffling** in **SequenceReader**.
- ▶ Added a new, internal test API.

- ▶ Added the ability to custom handle CPU affinity using the DALI\_AFFINITY\_MASK environment variable.

### Breaking API Changes

- ▶ None.

### Deprecated Features

- ▶ Dropped the Python 3.4 based **wh1** build.

### Known Issues

- ▶ The new video reader operator requires NVIDIA VIDEO CODEC SDK support in the platform. The NVIDIA GPU Cloud (NGC) optimized containers lack this functionality in the default configuration prior to the 19.01 version. To enable the functionality, run the container with the "video" capability enabled, as below:

```
-e "NVIDIA_DRIVER_CAPABILITIES=compute,utility,video"
```

- ▶ There is no clear distinction in the documentation between the operators supporting video sequences and the operators supporting images.

# Chapter 13.

## DALI RELEASE 0.6.1 BETA

The DALI 0.6.1 is a beta release. Hence, for all the features, the functionality and performance will likely be limited.

### Key Features and Enhancements

This DALI release includes the following key features and enhancements.

- ▶ Added the DALI TensorFlow plugin to improve out-of-the-box forward compatibility.
- ▶ Added a new package **nvidia-dali-tf-plugin** that allows user to compile DALI TensorFlow plugin using the locally installed version.
- ▶ Added the below examples:
  - ▶ VideoReader
  - ▶ PyTorch and MXNet examples with various readers
- ▶ Added a new operator variant **gpu box encoder**.
- ▶ Added a new JSON parser—now COCO reader start up time is significantly faster (~10x).
- ▶ Added the ability to automatically reset MXNet and PyTorch Python iterators.
- ▶ Enhanced the Video reader to link dynamically with the CUDA Driver API and **nvcuvid**.
- ▶ Added the ability to deliver exactly 1 epoch from **DALIGenericIterator** (PyTorch and MXNet).

### Breaking API Changes

- ▶ PyTorch iterator returns the exact number of samples per epoch, so the final batch could be smaller if epoch size is not divisible by the batch size. To keep the old behaviour when data is wrapped up, use **stop\_at\_epoch** argument.

## Using DALI 0.6.1 Beta

The 19.01 NVIDIA GPU Cloud (NGC) optimized container for MXNet, PyTorch, and TensorFlow includes an older version of DALI.

To upgrade to DALI 0.6.1 beta, follow the installation instructions in the [DALI Quick Start Guide](#).

Refer to the [DALI Developer Guide](#) for usage details.

## Known Issues

- ▶ The New Video reader operator requires NVIDIA VIDEO CODEC SDK support in the platform. The NVIDIA GPU Cloud (NGC) optimized containers lack this functionality in the default configuration prior to the 19.01 version, resulting in problem with loading `libnvcuvid.so`. To enable it, run the container with the 'video' capability enabled, as below:

```
-e "NVIDIA_DRIVER_CAPABILITIES=compute,utility,video"
```

# Chapter 14.

## DALI RELEASE 0.6 BETA

The DALI 0.6 is a beta release. Hence, for all the features, the functionality and performance will likely be limited.

### Key Features and Enhancements

This DALI release includes the following key features and enhancements.

- ▶ Added the CPU variant of DALI TensorFlow operator. Now the DALI pipeline can be scheduled on the CPU also.
- ▶ Added the TensorFlow operator API to allow defining an arbitrary number of outputs and their types.
- ▶ Added full DALI detection pipeline sample for SSD.
- ▶ Added clear mapping between operators and the supported device types in the API documentation.
- ▶ Created initial support for the video-based pipeline - [NVVL](#) functionality is available as `VideoReader`.
- ▶ Added the below new operators:
  - ▶ `BoxEncoder`
  - ▶ `VideoReader`
  - ▶ `SequenceCrop`
  - ▶ `SequenceReader`
- ▶ Added new variants of the below existing operators:
  - ▶ For the GPU:
    - ▶ `RandomBBoxCrop`
    - ▶ `BbFlip`
    - ▶ `Slice`
  - ▶ For the CPU:
    - ▶ `CropMirrorNormalize`, and

▶ **BBoxPaste**

### Breaking API Changes

- ▶ DALI TensorFlow operator has a new API - check the examples for the reference.
- ▶ PyTorch and MXNet python iterators API have changed - check the examples for the reference.

### Using DALI 0.6 Beta

The 19.01 NVIDIA GPU Cloud (NGC) optimized container for MXNet, PyTorch, and TensorFlow includes an older version of DALI. To upgrade to DALI 0.6 beta, follow the installation instructions in the [DALI Quick Start Guide](#).

Refer to the [DALI Developer Guide](#) for usage details.

### Known Issues

- ▶ The New Video reader operator requires NVIDIA VIDEO CODEC SDK support in the platform. The NVIDIA GPU Cloud (NGC) optimized containers lack this functionality in the default configuration prior to the 19.01 version, resulting in problem with loading `libnvcuvid.so`. To enable it, run the container with the 'video' capability enabled, as below:

```
-e "NVIDIA_DRIVER_CAPABILITIES=compute,utility,video"
```

# Chapter 15.

## DALI RELEASE 0.5 BETA

This DALI 0.5 release is a beta release.

### Key Features and Enhancements

This DALI release includes the following key features and enhancements.

- ▶ Fixed rare hangs during DALI pipeline destruction.
- ▶ Added ResNet50 example with TensorFlow.
- ▶ Added more verbose error messages to HostDecoder and nvJpegDecoder.
- ▶ Added pipelined execution completion callback setter.
- ▶ Added support for Tiff images.

### Breaking API Changes

- ▶ Random seed type changed from INT to INT64, therefore, serialized pipelines from versions prior to 0.5 are not compatible with the current DALI version.

### Using DALI 0.5 Beta

The 18.11 NVIDIA GPU Cloud (NGC) optimized container for MXNet, PyTorch, and TensorFlow, includes an older version of DALI. To upgrade to DALI 0.5 beta, follow the installation instructions in the [DALI Quick Start Guide](#).

Refer to the [DALI Developer Guide](#) for usage details.

### Known Issues

- ▶ This is a beta release. Hence, for all the features, the functionality and performance will likely be limited.

# Chapter 16.

## DALI RELEASE 0.4.1 BETA

This DALI 0.4.1 release is a beta release.

### Key Features and Enhancements

This DALI release includes the following key features and enhancements.

- ▶ Added TensorFlow 1.11 and 1.12 compatibility.
- ▶ Adjusted PyTorch example to use new nvJPEG API.
- ▶ Reduced DALI memory footprint.
- ▶ Fixed lack of ability to customize pipeline processing length.

### Using DALI 0.4.1 Beta

The 18.10 NVIDIA GPU Cloud (NGC) optimized container for MXNet, PyTorch, and TensorFlow, includes an older version of DALI. To upgrade to DALI 0.4.1 beta, follow the installation instructions in the [DALI Quick Start Guide](#).

Refer to the [DALI Developer Guide](#) for usage details.

### Known Issues

- ▶ The DALI integrated ResNet-50 samples in the 18.10 NGC TensorFlow and PyTorch containers may result in lower than expected performance results. We are working to address the issue in the next release.
- ▶ This is a beta release. All features are expected to be available, however, some aspects of functionality and performance will likely be limited compared to a non-beta release.



# Chapter 17.

## DALI RELEASE 0.4 BETA

This DALI 0.4 release is a beta release.

### Key Features and Enhancements

This DALI release includes the following key features and enhancements.

- ▶ Added basic operators for detection (SSD), such as COCO dataset reader (**COCOReader**), random crop operator with bounding boxes (**SSDRandomCrop**), and flip operator for bounding boxes (**BbFlip**).
- ▶ Added CPU versions of **Crop/CropCastPermute** operators.
- ▶ Added a random **Paste** operator.
- ▶ Upgraded OpenCV to 3.4.3.
- ▶ Upgraded MXNet to 1.3.0.
- ▶ Fixed parsing of JPEG headers by the Host Decoder known issue. The Host Decoder now handles all images from the ImageNet dataset.
- ▶ Added fine grained control over output buffers in the pipeline
- ▶ Updated to nvJPEG 0.2.0

### Breaking API Changes

The pipeline constructor signature has changed; an additional argument (**prefetch\_queue\_depth**) was added to allow defining depth of the prefetch queue at runtime.

## Using DALI 0.4 Beta

The 18.09 NVIDIA GPU Cloud (NGC) optimized container for MXNet, PyTorch, and TensorFlow, includes an older version of DALI. To upgrade to DALI 0.4 beta, follow the installation instructions in the [DALI Quick Start Guide](#).

Refer to the [DALI Developer Guide](#) for usage details.

## Deprecated Features

- ▶ DALI 0.4 is not compatible with TensorFlow 1.11. This will be addressed in the next release.

## Known Issues

- ▶ The DALI integrated ResNet-50 samples in the 18.10 NGC TensorFlow and PyTorch containers may result in lower than expected performance results. We are working to address the issue in the next release.
- ▶ This is a beta release. All features are expected to be available, however, some aspects of functionality and performance will likely be limited compared to a non-beta release.

# Chapter 18.

## DALI RELEASE 0.3 BETA

This DALI 0.3 release is a beta release.

### Key Features and Enhancements

This DALI release includes the following key features and enhancements.

- ▶ Updated PyTorch ResNet-50 example to obtain expected accuracy (Top1 76%).
- ▶ Introduced CPU variant of resize operator and added stand-alone flip operator.
- ▶ Added support for DALI to work with float16 data passed from Python.
- ▶ Added fallback to CPU for BMP images.
- ▶ Fixed training accuracy with TFRecord reader

### Using DALI 0.3 Beta

The 18.08 NVIDIA GPU Cloud (NGC) optimized container for MXNet, PyTorch, and TensorFlow, includes an older version of DALI. To upgrade to DALI 0.2 beta, follow the installation instructions in the [DALI Quick Start Guide](#).

Refer to the [DALI Developer Guide](#) for usage details.

### Known Issues

- ▶ HosDecoder cannot handle all jpeg files from ImageNet dataset. This will be fixed in the next release.
- ▶ On file systems where the directory entries are not stored in any order, the File Reader may assign different labels to the training and validation folders with the same name. This will be fixed in the next release.

- ▶ The DALI integrated ResNet-50 samples in the 18.09 NGC TensorFlow and PyTorch containers may result in lower than expected performance results. We are working to address the issue in the next release.
- ▶ This is a beta release. All features are expected to be available, however, some aspects of functionality and performance will likely be limited compared to a non-beta release.

# Chapter 19.

## DALI RELEASE 0.2 BETA

This DALI 0.2 release is a beta release.

### Key Features and Enhancements

This DALI release includes the following key features and enhancements.

- ▶ Added Sphinx based documentation that is in sync with the code on GitHub. For more information, see [DALI Master Branch User Guide](#).
- ▶ Build system has been refined, common errors have meaningful messages, improved localization and version detection of key dependency packages, such as, nvJPEG, libturbo-jpeg, and LMDB.
- ▶ Added **Unfused Crop** and **CropCastPermute** operators.
- ▶ Added improvements for TensorFlow plugin (polymorphism and shape argument for the output).
- ▶ Expanded examples of TensorFlow working with different readers, such as, MXNetReader, FileReader, and TFRecordReader.
- ▶ Updated nvJPEG to 0.1.4
- ▶ Added fallback to host decoder when image is not JPEG but PNG instead. For example, `n02105855_2933.JPEG` from ImageNet.

### Breaking API Changes

- ▶ The API for the **Resize** operator changed to match other similar operators like **ResizeCropMirror**.
- ▶ The API for the TensorFlow plugin changed to allow specifying the whole shape of the tensor instead of **N**, **H**, and **W** separately; which enables handling both **NCHW** and **NHWC** outputs.

- ▶ The type of labels produced by the TensorFlow plugin have changed. In DALI version 0.1.2, it was always `tf.float32`. In this release, a new optional parameter called `label_type` is introduced to the TensorFlow plugin to control the type of label. The default value for `label_type` is `tf.int64` to better align with the label type in TFRecord.

### Using DALI 0.2 Beta

The 18.08 NVIDIA GPU Cloud (NGC) optimized container for MXNet, PyTorch, and TensorFlow, includes an older version of DALI. To upgrade to DALI 0.2 beta, follow the installation instructions in the [DALI Quick Start Guide](#).

Refer to the [DALI Developer Guide](#) for usage details.

### Known Issues

- ▶ This is a beta release, therefore, not all functionality is fully supported and working. This beta release is meant for testing and research.
- ▶ The DALI integrated ResNet-50 samples in the 18.08 NGC TensorFlow and PyTorch containers have lower than expected accuracy and performance results. We are working to address the issue in the next release.
- ▶ This is a beta release. All features are expected to be available, however, some aspects of functionality and performance will likely be limited compared to a non-beta release.

# Chapter 20.

## DALI RELEASE 0.1.2 BETA

This DALI 0.1.2 release is a beta release.

### Key Features and Enhancements

This DALI release includes the following key features and enhancements.

- ▶ Fixed compatibility with TensorFlow 1.9.
- ▶ Updated to nvJPEG v0.1.2 to fix batched decoding when a batch contains both gray-scale and color images.
- ▶ Added Tensorflow 1.7 support.
- ▶ Improved overlap support when using DALI with multi-GPU in MXNet and PyTorch.

### Using DALI 0.1.2 Beta

The 18.07 optimized container for MXNet, PyTorch, and TensorFlow, includes an older version of DALI. To upgrade to DALI 0.2 beta, follow the installation instructions in the [DALI Quick Start Guide](#).

Refer to the [DALI Developer Guide](#) for usage details.

### Known Issues

- ▶ This is a beta release, therefore, not all functionality is fully supported and working. This beta release is meant for testing and research.
- ▶ This is a beta release. All features are expected to be available, however, some aspects of functionality and performance will likely be limited compared to a non-beta release.

# Chapter 21.

## DALI RELEASE 0.1.1 BETA

This is the first release of DALI. This DALI 0.1.1 release is a beta release.

### Key Features and Enhancements

This DALI release includes the following key features and enhancements.

#### Performance

On dense GPU systems, deep learning applications can be significantly bottlenecked on the CPU, limiting the overall performance and scalability of training and inference tasks. DALI enables offloading key deep learning augmentation steps on to GPUs, alleviating CPU bottleneck on the deep learning preprocessing pipelines. This results in out-of-box performance of overall training workflow and efficient utilization of multi-GPU resources on the system.

#### Drop-in Integration

DALI comes with built-in plugins for key frameworks such as MXNet, TensorFlow, and PyTorch. This enables automatic integration with frameworks so that researchers and developers can get up and running with DALI easily and quickly.

#### Flexibility

DALI supports multiple input data formats that are commonly used in computer vision deep learning applications, for example, JPEG images, raw formats, Lightning Memory-Mapped Database (LMDB), RecordIO and TFRecord. The flexibility of input data formats allows portability of training workflows across different frameworks and models, and helps to avoid intermediate data conversion steps. DALI enables better code reuse and maintainability with optimized building blocks and support for different data formats.

### Using DALI 0.1.1 Beta

Ensure you are familiar with the following notes when using this release.



- ▶ To install DALI, see the [DALI Quick Start Guide](#).



If you are using the 18.07 NGC optimized container for MXNet, PyTorch, or TensorFlow, you do not need to reinstall DALI. DALI now comes included in the container. Instead, start with the [Getting Started Tutorial](#).

- ▶ To interact with the code via GitHub, see the [Getting Started Tutorial](#).
- ▶ To learn how to define, build, and run a DALI pipeline, see the [DALI Developer Guide](#).

### Known Issues

- ▶ This is a beta release, therefore, not all functionality is fully supported and working. This beta release is meant for testing and research.
- ▶ This is a beta release. All features are expected to be available, however, some aspects of functionality and performance will likely be limited compared to a non-beta release.

## Notice

THE INFORMATION IN THIS GUIDE AND ALL OTHER INFORMATION CONTAINED IN NVIDIA DOCUMENTATION REFERENCED IN THIS GUIDE IS PROVIDED “AS IS.” NVIDIA MAKES NO WARRANTIES, EXPRESSED, IMPLIED, STATUTORY, OR OTHERWISE WITH RESPECT TO THE INFORMATION FOR THE PRODUCT, AND EXPRESSLY DISCLAIMS ALL IMPLIED WARRANTIES OF NONINFRINGEMENT, MERCHANTABILITY, AND FITNESS FOR A PARTICULAR PURPOSE. Notwithstanding any damages that customer might incur for any reason whatsoever, NVIDIA’s aggregate and cumulative liability towards customer for the product described in this guide shall be limited in accordance with the NVIDIA terms and conditions of sale for the product.

THE NVIDIA PRODUCT DESCRIBED IN THIS GUIDE IS NOT FAULT TOLERANT AND IS NOT DESIGNED, MANUFACTURED OR INTENDED FOR USE IN CONNECTION WITH THE DESIGN, CONSTRUCTION, MAINTENANCE, AND/OR OPERATION OF ANY SYSTEM WHERE THE USE OR A FAILURE OF SUCH SYSTEM COULD RESULT IN A SITUATION THAT THREATENS THE SAFETY OF HUMAN LIFE OR SEVERE PHYSICAL HARM OR PROPERTY DAMAGE (INCLUDING, FOR EXAMPLE, USE IN CONNECTION WITH ANY NUCLEAR, AVIONICS, LIFE SUPPORT OR OTHER LIFE CRITICAL APPLICATION). NVIDIA EXPRESSLY DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY OF FITNESS FOR SUCH HIGH RISK USES. NVIDIA SHALL NOT BE LIABLE TO CUSTOMER OR ANY THIRD PARTY, IN WHOLE OR IN PART, FOR ANY CLAIMS OR DAMAGES ARISING FROM SUCH HIGH RISK USES.

NVIDIA makes no representation or warranty that the product described in this guide will be suitable for any specified use without further testing or modification. Testing of all parameters of each product is not necessarily performed by NVIDIA. It is customer’s sole responsibility to ensure the product is suitable and fit for the application planned by customer and to do 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 guide. NVIDIA does not accept any 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 guide, or (ii) customer product designs.

Other than the right for customer to use the information in this guide with the product, no other license, either expressed or implied, is hereby granted by NVIDIA under this guide. Reproduction of information in this guide is permissible only if reproduction is approved by NVIDIA in writing, is reproduced without alteration, and is accompanied by all associated conditions, limitations, and notices.

## Trademarks

NVIDIA, the NVIDIA logo, and cuBLAS, CUDA, cuDNN, DALI, DIGITS, DGX, DGX-1, DGX-2, DGX Station, DLProf, Jetson, Kepler, Maxwell, NCCL, Nsight Compute, Nsight Systems, NvCaffe, PerfWorks, Pascal, SDK Manager, Tegra, TensorRT, TensorRT Inference Server, Tesla, TF-TRT, 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

© 2019 NVIDIA Corporation. All rights reserved.

[www.nvidia.com](http://www.nvidia.com)

