



TORCH

RN-08469-001_v18.08 | January 2020

Release Notes



TABLE OF CONTENTS

Chapter 1. Torch Overview.....	1
Chapter 2. Pulling A Container.....	2
Chapter 3. Running Torch.....	3
Chapter 4. Torch Release 18.08.....	5
Chapter 5. Torch Release 18.07.....	7
Chapter 6. Torch Release 18.06.....	9
Chapter 7. Torch Release 18.05.....	11
Chapter 8. Torch Release 18.04.....	13
Chapter 9. Torch Release 18.03.....	15
Chapter 10. Torch Release 18.02.....	17
Chapter 11. Torch Release 18.01.....	19
Chapter 12. Torch Release 17.12.....	21
Chapter 13. Torch Release 17.11.....	23
Chapter 14. Torch Release 17.10.....	25
Chapter 15. Torch Release 17.09.....	27
Chapter 16. Torch Release 17.07.....	29
Chapter 17. Torch Release 17.06.....	30
Chapter 18. Torch Release 17.05.....	32
Chapter 19. Torch Release 17.04.....	34
Chapter 20. Torch Release 17.03.....	35
Chapter 21. Torch Release 17.02.....	36
Chapter 22. Torch Release 17.01.....	37
Chapter 23. Torch Release 16.12.....	38

Chapter 1.

TORCH OVERVIEW

The NVIDIA Deep Learning SDK accelerates widely-used deep learning frameworks such as [Torch™](#).

Torch is a scientific computing framework with wide support for deep learning algorithms. Thanks to an easy and fast scripting language, Lua, and an underlying C/CUDA® implementation, Torch is easy to use and is efficient.

Torch offers popular neural network and optimization libraries that are easy to use yet provide maximum flexibility to build complex neural network topologies.

See `/workspace/README.md` inside the container for information on customizing your Torch image. For more information about Torch, see:

- ▶ [Torch website](#)
- ▶ [Torch project](#)

This document describes the key features, software enhancements and improvements, any known issues, and how to run this container.

Chapter 2.

PULLING A CONTAINER

Before you can pull a container from the NGC container registry, you must have Docker installed. For DGX users, this is explained in [Preparing to use NVIDIA Containers Getting Started Guide](#).

For users other than DGX, follow the NVIDIA® GPU Cloud™ (NGC) container registry [installation documentation](#) based on your platform.

You must also have access and logged into the NGC container registry as explained in the [NGC Getting Started Guide](#).

There are four repositories where you can find the NGC docker containers.

`nvcr.io/nvidia`

The deep learning framework containers are stored in the **`nvcr.io/nvidia`** repository.

`nvcr.io/hpc`

The HPC containers are stored in the **`nvcr.io/hpc`** repository.

`nvcr.io/nvidia-hpcvis`

The HPC visualization containers are stored in the **`nvcr.io/nvidia-hpcvis`** repository.

`nvcr.io/partner`

The partner containers are stored in the **`nvcr.io/partner`** repository. Currently the partner containers are focused on Deep Learning or Machine Learning, but that doesn't mean they are limited to those types of containers.

Chapter 3.

RUNNING TORCH

Before you can run an NGC deep learning framework container, your Docker environment must support NVIDIA GPUs. To run a container, issue the appropriate command as explained in the [Running A Container](#) chapter in the *NVIDIA Containers And Frameworks User Guide* and specify the registry, repository, and tags.

On a system with GPU support for NGC containers, the following occurs when running a container:

- ▶ The Docker engine loads the image into a container which runs the software.
- ▶ You define the runtime resources of the container by including additional flags and settings that are used with the command. These flags and settings are described in [Running A Container](#).
- ▶ The GPUs are explicitly defined for the Docker container (defaults to all GPUs, can be specified using `NV_GPU` environment variable).

The method implemented in your system depends on the DGX OS version installed (for DGX systems), the specific NGC Cloud Image provided by a Cloud Service Provider, or the software that you have installed in preparation for running NGC containers on TITAN PCs, Quadro PCs, or vGPUs.

1. Issue the command for the applicable release of the container that you want. The following command assumes you want to pull the latest container.

```
docker pull nvcr.io/nvidia/torch:18.08
```

2. Open a command prompt and paste the pull command. The pulling of the container image begins. Ensure the pull completes successfully before proceeding to the next step.
3. Run the container image. To run the container, choose interactive mode or non-interactive mode.

- a) **Interactive mode:** Open a command prompt and issue:

```
docker run --gpus all -it --rm -v local_dir:container_dir  
nvcr.io/nvidia/torch:<xx.xx>
```

- b) **Non-interactive mode:** Open a command prompt and issue:

```
docker run --gpus all --rm -v local_dir:container_dir  
nvcv.io/nvidia/torch:<xx.xx> <command>
```

You might want to pull in data and model descriptions from locations outside the container for use by Torch or save results to locations outside the container. To accomplish this, the easiest method is to mount one or more host directories as [Docker data volumes](#).



Deep Learning GPU Training System™ (DIGITS) uses shared memory to share data between processes. For example, if you use Torch multiprocessing for multi-threaded data loaders, the default shared memory segment size that the container runs with may not be enough. Therefore, you should increase the shared memory size by issuing either:

```
--ipc=host
```

or

```
--shm-size=<requested memory size>
```

in the command line to:

```
docker run --gpus all
```

Chapter 4.

TORCH RELEASE 18.08

The NVIDIA container image of Torch, release 18.08, is available.

Contents of Torch

This container image contains the complete source of the version of NVIDIA Torch in `/opt/torch`. It is pre-built and installed into the `/usr/local/[bin,share,lib]` directories in the container image.

The NVIDIA[®] Collective Communications Library[™] (NCCL) (NCCL) library and Torch bindings for NCCL are installed in this container, and models using `DataParallelTable` can easily leverage this library for fast parallel training.

The container also includes the following:

- ▶ [Ubuntu 16.04](#) including [Python 2.7](#) environment
- ▶ [NVIDIA CUDA 9.0.176](#) (see [Errata](#) section and 2.1) including [CUDA[®] Basic Linear Algebra Subroutines library[™] \(cuBLAS\) 9.0.425](#)
- ▶ [NVIDIA CUDA[®] Deep Neural Network library[™] \(cuDNN\) 7.2.1](#)
- ▶ [NCCL 2.2.13](#) (optimized for [NVLink[™]](#))

Driver Requirements

Release 18.08 is based on CUDA 9, which requires [NVIDIA Driver](#) release 384.xx.

Key Features and Enhancements

This Torch release includes the following key features and enhancements.

- ▶ [Torch](#) container image version 18.08 is based on [Torch7](#).
- ▶ Latest version of [cuDNN 7.2.1](#).
- ▶ [Ubuntu 16.04](#) with July 2018 updates

Announcements

We are continuing to incorporate monthly Torch upstream changes as well as upgrade NVIDIA libraries, such as cuDNN, cuBLAS, NCCL, and the Ubuntu OS. However, we will be discontinuing container updates once the next major CUDA version is released.

Known Issues

There are no known issues in this release.

Chapter 5.

TORCH RELEASE 18.07

The NVIDIA container image of Torch, release 18.07, is available.

Contents of Torch

This container image contains the complete source of the version of NVIDIA Torch in `/opt/torch`. It is pre-built and installed into the `/usr/local/[bin,share,lib]` directories in the container image.

The NVIDIA[®] Collective Communications Library[™] (NCCL) (NCCL) library and Torch bindings for NCCL are installed in this container, and models using `DataParallelTable` can easily leverage this library for fast parallel training.

The container also includes the following:

- ▶ Ubuntu 16.04 including Python 2.7 environment
- ▶ NVIDIA CUDA 9.0.176 (see Errata section and 2.1) including CUDA[®] Basic Linear Algebra Subroutines library[™] (cuBLAS) 9.0.425
- ▶ NVIDIA CUDA[®] Deep Neural Network library[™] (cuDNN) 7.1.4
- ▶ NCCL 2.2.13 (optimized for NVLink[™])

Driver Requirements

Release 18.07 is based on CUDA 9, which requires [NVIDIA Driver](#) release 384.xx.

Key Features and Enhancements

This Torch release includes the following key features and enhancements.

- ▶ Torch container image version 18.07 is based on Torch7.
- ▶ Latest version of CUDA[®] Basic Linear Algebra Subroutines library[™] (cuBLAS) 9.0.425.
- ▶ Ubuntu 16.04 with June 2018 updates

Announcements

We are continuing to incorporate monthly Torch upstream changes as well as upgrade NVIDIA libraries, such as cuDNN, cuBLAS, NCCL, and the Ubuntu OS. However, we will be discontinuing container updates once the next major CUDA version is released.

Known Issues

There are no known issues in this release.

Chapter 6.

TORCH RELEASE 18.06

The NVIDIA container image of Torch, release 18.06, is available.

Contents of Torch

This container image contains the complete source of the version of NVIDIA Torch in `/opt/torch`. It is pre-built and installed into the `/usr/local/[bin,share,lib]` directories in the container image.

The NVIDIA[®] Collective Communications Library[™] (NCCL) (NCCL) library and Torch bindings for NCCL are installed in this container, and models using `DataParallelTable` can easily leverage this library for fast parallel training.

The container also includes the following:

- ▶ [Ubuntu 16.04](#) including [Python 2.7](#) environment
- ▶ [NVIDIA CUDA 9.0.176](#) (see [Errata](#) section and 2.1) including [CUDA[®] Basic Linear Algebra Subroutines library[™] \(cuBLAS\) 9.0.333](#) (see section 2.3.1)
- ▶ [NVIDIA CUDA[®] Deep Neural Network library[™] \(cuDNN\) 7.1.4](#)
- ▶ [NCCL 2.2.13](#) (optimized for [NVLink[™]](#))

Driver Requirements

Release 18.06 is based on CUDA 9, which requires [NVIDIA Driver](#) release 384.xx.

Key Features and Enhancements

This Torch release includes the following key features and enhancements.

- ▶ [Torch](#) container image version 18.06 is based on [Torch7](#).
- ▶ [Ubuntu 16.04](#) with May 2018 updates

Announcements

We are continuing to incorporate monthly Torch upstream changes as well as upgrade NVIDIA libraries, such as cuDNN, cuBLAS, NCCL, and the Ubuntu OS. However, we will be discontinuing container updates once the next major CUDA version is released.

Known Issues

There are no known issues in this release.

Chapter 7.

TORCH RELEASE 18.05

The NVIDIA container image of Torch, release 18.05, is available.

Contents of Torch

This container image contains the complete source of the version of NVIDIA Torch in `/opt/torch`. It is pre-built and installed into the `/usr/local/[bin,share,lib]` directories in the container image.

The NVIDIA[®] Collective Communications Library[™] (NCCL) (NCCL) library and Torch bindings for NCCL are installed in this container, and models using `DataParallelTable` can easily leverage this library for fast parallel training.

The container also includes the following:

- ▶ [Ubuntu 16.04](#) including [Python 2.7](#) environment
- ▶ [NVIDIA CUDA 9.0.176](#) (see [Errata](#) section and 2.1) including [CUDA[®] Basic Linear Algebra Subroutines library[™] \(cuBLAS\) 9.0.333](#) (see section 2.3.1)
- ▶ [NVIDIA CUDA[®] Deep Neural Network library[™] \(cuDNN\) 7.1.2](#)
- ▶ [NCCL 2.1.15](#) (optimized for [NVLink[™]](#))

Driver Requirements

Release 18.05 is based on CUDA 9, which requires [NVIDIA Driver](#) release 384.xx.

Key Features and Enhancements

This Torch release includes the following key features and enhancements.

- ▶ [Torch](#) container image version 18.05 is based on [Torch7](#).
- ▶ [Ubuntu 16.04](#) with April 2018 updates

Announcements

We are continuing to incorporate monthly Torch upstream changes as well as upgrade NVIDIA libraries, such as cuDNN, cuBLAS, NCCL, and the Ubuntu OS. However, we will be discontinuing container updates once the next major CUDA version is released.

Known Issues

There are no known issues in this release.

Chapter 8.

TORCH RELEASE 18.04

The NVIDIA container image of Torch, release 18.04, is available.

Contents of Torch

This container image contains the complete source of the version of NVIDIA Torch in `/opt/torch`. It is pre-built and installed into the `/usr/local/[bin,share,lib]` directories in the container image.

The NVIDIA[®] Collective Communications Library[™] (NCCL) (NCCL) library and Torch bindings for NCCL are installed in this container, and models using `DataParallelTable` can easily leverage this library for fast parallel training.

The container also includes the following:

- ▶ [Ubuntu 16.04](#) including [Python 2.7](#) environment
- ▶ [NVIDIA CUDA 9.0.176](#) (see [Errata](#) section and 2.1) including [CUDA[®] Basic Linear Algebra Subroutines library[™] \(cuBLAS\) 9.0.333](#) (see section 2.3.1)
- ▶ [NVIDIA CUDA[®] Deep Neural Network library[™] \(cuDNN\) 7.1.1](#)
- ▶ [NCCL 2.1.15](#) (optimized for [NVLink[™]](#))

Driver Requirements

Release 18.04 is based on CUDA 9, which requires [NVIDIA Driver](#) release 384.xx.

Key Features and Enhancements

This Torch release includes the following key features and enhancements.

- ▶ [Torch](#) container image version 18.04 is based on [Torch7](#).
- ▶ Latest version of [NCCL 2.1.15](#)
- ▶ [Ubuntu 16.04](#) with March 2018 updates

Announcements

We are continuing to incorporate monthly Torch upstream changes as well as upgrade NVIDIA libraries, such as cuDNN, cuBLAS, NCCL, and the Ubuntu OS. However, we will be discontinuing container updates once the next major CUDA version is released.

Known Issues

There are no known issues in this release.

Chapter 9.

TORCH RELEASE 18.03

The NVIDIA container image of Torch, release 18.03, is available.

Contents of Torch

This container image contains the complete source of the version of NVIDIA Torch in `/opt/torch`. It is pre-built and installed into the `/usr/local/[bin,share,lib]` directories in the container image.

The NVIDIA[®] Collective Communications Library[™] (NCCL) (NCCL) library and Torch bindings for NCCL are installed in this container, and models using `DataParallelTable` can easily leverage this library for fast parallel training.

The container also includes the following:

- ▶ Ubuntu 16.04 including Python 2.7 environment
- ▶ NVIDIA CUDA 9.0.176 (see Errata section and 2.1) including CUDA[®] Basic Linear Algebra Subroutines library[™] (cuBLAS) 9.0.333 (see section 2.3.1)
- ▶ NVIDIA CUDA[®] Deep Neural Network library[™] (cuDNN) 7.1.1
- ▶ NCCL 2.1.2 (optimized for NVLink[™])

Driver Requirements

Release 18.03 is based on CUDA 9, which requires [NVIDIA Driver](#) release 384.xx.

Key Features and Enhancements

This Torch release includes the following key features and enhancements.

- ▶ Torch container image version 18.03 is based on Torch7.
- ▶ Latest version of cuBLAS 9.0.333
- ▶ Latest version of cuDNN 7.1.1
- ▶ Ubuntu 16.04 with February 2018 updates

Announcements

Due to minimal upstream activity, we will no longer incorporate upstream changes going forward. We will continue to upgrade NVIDIA libraries, such as cuDNN, cuBLAS, and NCCL, as well as the Ubuntu OS for another half a year until we discontinue support.

Known Issues

There are no known issues in this release.

Chapter 10.

TORCH RELEASE 18.02

The NVIDIA container image of Torch, release 18.02, is available.

Torch container image version 18.02 is based on Torch7.

Contents of Torch

This container image contains the complete source of the version of NVIDIA Torch in `/opt/torch`. It is pre-built and installed into the `/usr/local/[bin,share,lib]` directories in the container image.

The NVIDIA[®] Collective Communications Library[™] (NCCL) (NCCL) library and Torch bindings for NCCL are installed in this container, and models using `DataParallelTable` can easily leverage this library for fast parallel training.

The container also includes the following:

- ▶ Ubuntu 16.04 including Python 2.7 environment
- ▶ NVIDIA CUDA 9.0.176 including:
 - ▶ CUDA[®] Basic Linear Algebra Subroutines library[™] (cuBLAS) 9.0.282 Patch 2 which is installed by default
 - ▶ cuBLAS 9.0.234 Patch 1 as a debian file. Installing Patch 1 by issuing the `dpkg -i /opt/cuda-cublas-9-0_9.0.234-1_amd64.deb` command is the workaround for the known issue described below.
- ▶ NVIDIA CUDA[®] Deep Neural Network library[™] (cuDNN) 7.0.5
- ▶ NCCL 2.1.2 (optimized for NVLink[™])

Driver Requirements

Release 18.02 is based on CUDA 9, which requires NVIDIA Driver release 384.xx.

Key Features and Enhancements

This Torch release includes the following key features and enhancements.

- ▶ Incorporated all upstream changes.



We will no longer incorporate upstream changes after the 18.02 release due to the lack of upstream activity.

- ▶ Latest version of cuBLAS
- ▶ Ubuntu 16.04 with January 2018 updates

Known Issues

cuBLAS 9.0.282 regresses RNN seq2seq FP16 performance for a small subset of input sizes. This issue should be fixed in the next update. As a workaround, install cuBLAS 9.0.234 Patch 1 by issuing the `dpkg -i /opt/cuda-cublas-9-0_9.0.234-1_amd64.deb` command.

Chapter 11.

TORCH RELEASE 18.01

The NVIDIA container image of Torch, release 18.01, is available.

Torch container image version 18.01 is based on Torch7.

Contents of Torch

This container image contains the complete source of the version of NVIDIA Torch in `/opt/torch`. It is pre-built and installed into the `/usr/local/[bin,share,lib]` directories in the container image.

The NVIDIA[®] Collective Communications Library[™] (NCCL) (NCCL) library and Torch bindings for NCCL are installed in this container, and models using `DataParallelTable` can easily leverage this library for fast parallel training.

The container also includes the following:

- ▶ Ubuntu 16.04 including Python 2.7 environment
- ▶ NVIDIA CUDA 9.0.176 including CUDA[®] Basic Linear Algebra Subroutines library[™] (cuBLAS) 9.0.282
- ▶ NVIDIA CUDA[®] Deep Neural Network library[™] (cuDNN) 7.0.5
- ▶ NCCL 2.1.2 (optimized for NVLink[™])

Driver Requirements

Release 18.01 is based on CUDA 9, which requires NVIDIA Driver release 384.xx.

Key Features and Enhancements

This Torch release includes the following key features and enhancements.

- ▶ Latest version of cuBLAS
- ▶ Latest version of cuDNN
- ▶ Latest version of NCCL
- ▶ Ubuntu 16.04 with December 2017 updates

Known Issues

cuBLAS 9.0.282 regresses RNN seq2seq FP16 performance for a small subset of input sizes. As a workaround, revert back to the 11.12 container.

Chapter 12.

TORCH RELEASE 17.12

The NVIDIA container image of Torch, release 17.12, is available.

Torch container image version 17.12 is based on Torch7.

Contents of Torch

This container image contains the complete source of the version of NVIDIA Torch in `/opt/torch`. It is pre-built and installed into the `/usr/local/[bin,share,lib]` directories in the container image.

The NVIDIA[®] Collective Communications Library[™] (NCCL) (NCCL) library and Torch bindings for NCCL are installed in this container, and models using `DataParallelTable` can easily leverage this library for fast parallel training.

The container also includes the following:

- ▶ Ubuntu 16.04
- ▶ NVIDIA CUDA 9.0.176 including CUDA[®] Basic Linear Algebra Subroutines library[™] (cuBLAS) 9.0.234
- ▶ NVIDIA CUDA[®] Deep Neural Network library[™] (cuDNN) 7.0.5
- ▶ NCCL 2.1.2 (optimized for NVLink[™])

Driver Requirements

Release 17.12 is based on CUDA 9, which requires NVIDIA Driver release 384.xx.

Key Features and Enhancements

This Torch release includes the following key features and enhancements.

- ▶ Latest version of CUDA
- ▶ Latest version of cuDNN
- ▶ Latest version of NCCL
- ▶ Ubuntu 16.04 with November 2017 updates

Known Issues

There are no known issues in this release.

Chapter 13.

TORCH RELEASE 17.11

The NVIDIA container image of Torch, release 17.11, is available.

Torch container image version 17.11 is based on Torch7.

Contents of Torch

This container image contains the complete source of the version of NVIDIA Torch in `/opt/torch`. It is pre-built and installed into the `/usr/local/[bin,share,lib]` directories in the container image.

The NVIDIA[®] Collective Communications Library[™] (NCCL) (NCCL) library and Torch bindings for NCCL are installed in this container, and models using `DataParallelTable` can easily leverage this library for fast parallel training.

The container also includes the following:

- ▶ Ubuntu 16.04
- ▶ NVIDIA CUDA 9.0.176 including CUDA[®] Basic Linear Algebra Subroutines library[™] (cuBLAS) 9.0.234
- ▶ NVIDIA CUDA[®] Deep Neural Network library[™] (cuDNN) 7.0.4
- ▶ NCCL 2.1.2 (optimized for NVLink[™])

Driver Requirements

Release 17.11 is based on CUDA 9, which requires NVIDIA Driver release 384.xx.

Key Features and Enhancements

This Torch release includes the following key features and enhancements.

- ▶ Latest version of CUDA
- ▶ Latest version of cuDNN
- ▶ Latest version of NCCL
- ▶ Ubuntu 16.04 with October 2017 updates

Known Issues

There are no known issues in this release.

Chapter 14.

TORCH RELEASE 17.10

The NVIDIA container image of Torch, release 17.10, is available.

[Torch](#) container image version 17.10 is based on [Torch7](#).

Contents of Torch

This container image contains the complete source of the version of NVIDIA Torch in `/opt/torch`. It is pre-built and installed into the `/usr/local/[bin,share,lib]` directories in the container image.

The NVIDIA[®] Collective Communications Library[™] (NCCL) (NCCL) library and Torch bindings for NCCL are installed in this container, and models using `DataParallelTable` can easily leverage this library for fast parallel training.

The container also includes the following:

- ▶ [Ubuntu 16.04](#)
- ▶ [NVIDIA CUDA 9.0](#)
- ▶ [NVIDIA CUDA[®] Deep Neural Network library[™] \(cuDNN\) 7.0.3](#)
- ▶ [NCCL 2.0.5 \(optimized for NVLink[™]\)](#)

Driver Requirements

Release 17.10 is based on CUDA 9, which requires [NVIDIA Driver](#) release 384.xx.

Key Features and Enhancements

This Torch release includes the following key features and enhancements.

- ▶ Latest version of CUDA
- ▶ Latest version of cuDNN
- ▶ Latest version of NCCL
- ▶ Ubuntu 16.04 with September 2017 updates

Known Issues

There are no known issues in this release.

Chapter 15.

TORCH RELEASE 17.09

The NVIDIA container image of Torch, release 17.09, is available.

[Torch](#) container image version 17.09 is based on [Torch7](#).

Contents of Torch

This container image contains the complete source of the version of NVIDIA Torch in `/opt/torch`. It is pre-built and installed into the `/usr/local/[bin,share,lib]` directories in the container image.

The NVIDIA[®] Collective Communications Library[™] (NCCL) (NCCL) library and Torch bindings for NCCL are installed in this container, and models using `DataParallelTable` can easily leverage this library for fast parallel training.

The container also includes the following:

- ▶ [Ubuntu 16.04](#)
- ▶ [NVIDIA CUDA 9.0](#)
- ▶ [NVIDIA CUDA[®] Deep Neural Network library[™] \(cuDNN\) 7.0.2](#)
- ▶ [NCCL 2.0.5 \(optimized for NVLink[™]\)](#)

Driver Requirements

Release 17.09 is based on CUDA 9, which requires [NVIDIA Driver](#) release 384.xx.

Key Features and Enhancements

This Torch release includes the following key features and enhancements.

- ▶ Matrix multiplication operations on FP16 inputs use Tensor Core Ops when available
- ▶ Latest version of CUDA
- ▶ Latest version of cuDNN
- ▶ Latest version of NCCL

- ▶ Ubuntu 16.04 with August 2017 updates

Known Issues

There are no known issues in this release.

Chapter 16.

TORCH RELEASE 17.07

The NVIDIA container image of Torch, release 17.07, is available.

Torch container image version 17.07 is based on Torch7.

Contents of Torch

This container image contains the complete source of the version of NVIDIA Torch in `/opt/torch`. It is pre-built and installed into the `/usr/local/[bin,share,lib]` directories in the container image.

The NVIDIA[®] Collective Communications Library[™] (NCCL) (NCCL) library and Torch bindings for NCCL are installed in this container, and models using `DataParallelTable` can easily leverage this library for fast parallel training.

The container also includes the following:

- ▶ Ubuntu 16.04
- ▶ NVIDIA CUDA 8.0.61.2 including CUDA[®] Basic Linear Algebra Subroutines library[™] (cuBLAS) Patch 2
- ▶ NVIDIA CUDA[®] Deep Neural Network library[™] (cuDNN) 6.0.21
- ▶ NVIDIA NCCL 2.0.3 (optimized for NVLink[™])

Key Features and Enhancements

This Torch release includes the following key features and enhancements.

- ▶ Ubuntu 16.04 with June 2017 updates

Known Issues

There are no known issues in this release.

Chapter 17.

TORCH RELEASE 17.06

The NVIDIA container image of Torch, release 17.06, is available.

Torch container image version 17.06 is based on Torch7.

Contents of Torch

This container image contains the complete source of the version of NVIDIA Torch in `/opt/torch`. It is pre-built and installed into the `/usr/local/[bin,share,lib]` directories in the container image.

The NVIDIA[®] Collective Communications Library[™] (NCCL) (NCCL) library and Torch bindings for NCCL are installed in this container, and models using `DataParallelTable` can easily leverage this library for fast parallel training.

The container also includes the following:

- ▶ Ubuntu 16.04
- ▶ NVIDIA CUDA 8.0.61
- ▶ NVIDIA CUDA[®] Deep Neural Network library[™] (cuDNN) 6.0.21
- ▶ NVIDIA NCCL 1.6.1 (optimized for NVLink[™])

Key Features and Enhancements

This Torch release includes the following key features and enhancements.

- ▶ Ubuntu 16.04 with May 2017 updates

Known Issues

The NCCL library version 1.6.1 included in this image, modifies the output buffers on all GPUs during in-place `ncclReduce()` operations, whereas normally only the "root" (target) device's output buffer should be modified. This is fixed in later versions of NCCL, as will be packaged in later versions of this image. As a workaround, either use `ncclAllReduce()`, which correctly modifies output buffers of all GPUs to the same

values, or use out-of-place `ncclReduce()`, wherein the output buffer is distinct from the input buffer.

Chapter 18.

TORCH RELEASE 17.05

The NVIDIA container image of Torch, release 17.05, is available.

Torch container image version 17.05 is based on Torch7.

Contents of Torch

This container image contains the complete source of the version of NVIDIA Torch in `/opt/torch`. It is pre-built and installed into the `/usr/local/[bin,share,lib]` directories in the container image.

The NVIDIA[®] Collective Communications Library[™] (NCCL) (NCCL) library and Torch bindings for NCCL are installed in this container, and models using `DataParallelTable` can easily leverage this library for fast parallel training.

The container also includes the following:

- ▶ Ubuntu 16.04
- ▶ NVIDIA CUDA 8.0.61
- ▶ NVIDIA CUDA[®] Deep Neural Network library[™] (cuDNN) 6.0.21
- ▶ NVIDIA NCCL 1.6.1 (optimized for NVLink[™])

Key Features and Enhancements

This Torch release includes the following key features and enhancements.

- ▶ Latest cuDNN release
- ▶ Ubuntu 16.04 with April 2017 updates

Known Issues

The NCCL library version 1.6.1 included in this image, modifies the output buffers on all GPUs during in-place `ncclReduce()` operations, whereas normally only the "root" (target) device's output buffer should be modified. This is fixed in later versions of NCCL, as will be packaged in later versions of this image. As a workaround, either use

`ncclAllReduce()`, which correctly modifies output buffers of all GPUs to the same values, or use out-of-place `ncclReduce()`, wherein the output buffer is distinct from the input buffer.

Chapter 19.

TORCH RELEASE 17.04

The NVIDIA container image of Torch, release 17.04, is available.

[Torch](#) container image version 17.04 is based on [Torch7](#).

Contents of Torch

This container image contains the complete source of the version of NVIDIA Torch in `/opt/torch`. It is pre-built and installed into the `/usr/local/[bin,share,lib]` directories in the container image.

The NVIDIA[®] Collective Communications Library[™] (NCCL) (NCCL) library and Torch bindings for NCCL are installed in this container, and models using `DataParallelTable` can easily leverage this library for fast parallel training.

The container also includes the following:

- ▶ [Ubuntu 16.04](#)
- ▶ [NVIDIA CUDA 8.0.61](#)
- ▶ [NVIDIA CUDA[®] Deep Neural Network library[™] \(cuDNN\) 6.0.20](#)
- ▶ [NVIDIA NCCL 1.6.1 \(optimized for NVLink[™]\)](#)

Key Features and Enhancements

This Torch release includes the following key features and enhancements.

- ▶ Synced to upstream Torch version as of March 2017
- ▶ Ubuntu 16.04 with March 2017 updates

Known Issues

There are no known issues in this release.

Chapter 20.

TORCH RELEASE 17.03

The NVIDIA container image of Torch, release 17.03, is available.

[Torch](#) container image version 17.03 is based on [Torch7](#).

Contents of Torch

This container image contains the complete source of the version of NVIDIA Torch in `/opt/torch`. It is pre-built and installed into the `/usr/local/[bin,share,lib]` directories in the container image.

The NVIDIA[®] Collective Communications Library[™] (NCCL) (NCCL) library and Torch bindings for NCCL are installed in this container, and models using `DataParallelTable` can easily leverage this library for fast parallel training.

The container also includes the following:

- ▶ [Ubuntu 16.04](#)
- ▶ [NVIDIA CUDA 8.0.61](#)
- ▶ [NVIDIA CUDA[®] Deep Neural Network library[™] \(cuDNN\) 6.0.20](#)
- ▶ [NVIDIA NCCL 1.6.1 \(optimized for NVLink[™]\)](#)

Key Features and Enhancements

This Torch release includes the following key features and enhancements.

- ▶ [Ubuntu 16.04 with February 2017 updates](#)

Known Issues

There are no known issues in this release.

Chapter 21.

TORCH RELEASE 17.02

The NVIDIA container image of Torch, release 17.02, is available.

Torch container image version 17.02 is based on Torch7.

Contents of Torch

This container image contains the complete source of the version of NVIDIA Torch in `/opt/torch`. It is pre-built and installed into the `/usr/local/[bin,share,lib]` directories in the container image.

The NVIDIA[®] Collective Communications Library[™] (NCCL) (NCCL) library and Torch bindings for NCCL are installed in this container, and models using `DataParallelTable` can easily leverage this library for fast parallel training.

The container also includes the following:

- ▶ Ubuntu 14.04
- ▶ NVIDIA CUDA 8.0.61
- ▶ NVIDIA CUDA[®] Deep Neural Network library[™] (cuDNN) 6.0.13
- ▶ NVIDIA NCCL 1.6.1 (optimized for NVLink[™])

Key Features and Enhancements

This Torch release includes the following key features and enhancements.

- ▶ Pseudo-FP16 support in cuDNN layers
- ▶ Improved performance of cat operation
- ▶ Added support for half tensor on the host side
- ▶ Ubuntu 14.04 with January 2017 updates

Known Issues

There are no known issues in this release.

Chapter 22.

TORCH RELEASE 17.01

The NVIDIA container image of Torch, release 17.01, is available.

Torch container image version 17.01 is based on Torch7.

Contents of Torch

This container image contains the complete source of the version of NVIDIA Torch in `/opt/torch`. It is pre-built and installed into the `/usr/local/[bin,share,lib]` directories in the container image.

The NVIDIA[®] Collective Communications Library[™] (NCCL) (NCCL) library and Torch bindings for NCCL are installed in this container, and models using `DataParallelTable` can easily leverage this library for fast parallel training.

The container also includes the following:

- ▶ Ubuntu 14.04
- ▶ NVIDIA CUDA 8.0.54
- ▶ NVIDIA CUDA[®] Deep Neural Network library[™] (cuDNN) 6.0.10
- ▶ NVIDIA NCCL 1.6.1 (optimized for NVLink[™])

Key Features and Enhancements

This Torch release includes the following key features and enhancements.

- ▶ Ubuntu 14.04 with December 2016 updates

Known Issues

There are no known issues in this release.

Chapter 23.

TORCH RELEASE 16.12

The NVIDIA container image of Torch, release 16.12, is available.

Torch container image version 16.12 is based on Torch7.

Contents of Torch

This container image contains the complete source of the version of NVIDIA Torch in `/opt/torch`. It is pre-built and installed into the `/usr/local/[bin,share,lib]` directories in the container image.

The NVIDIA[®] Collective Communications Library[™] (NCCL) (NCCL) library and Torch bindings for NCCL are installed in this container, and models using `DataParallelTable` can easily leverage this library for fast parallel training.

The container also includes the following:

- ▶ Ubuntu 14.04
- ▶ NVIDIA CUDA 8.0.54
- ▶ NVIDIA CUDA[®] Deep Neural Network library[™] (cuDNN) 6.0.5
- ▶ NVIDIA NCCL 1.6.1 (optimized for NVLink[™])

Key Features and Enhancements

This Torch release includes the following key features and enhancements.

- ▶ Supports FP32 and FP16 storage and FP32 arithmetic
 - ▶ Pseudo-FP16 support in cuDNN layers
- ▶ Optimized multi-GPU training
 - ▶ Seamless NCCL integration (opt-in)
 - ▶ Better control over workspace memory usage via support the new `cudaDnnFindEx()` routine; enable in model script by adding `cudaDnn.useFindEx = true`

- ▶ Supports recurrent neural networks (RNNs)
 - ▶ Supports cuDNN RNN layers
 - ▶ Tuned RNN performance
- ▶ Lua frontend
- ▶ Ubuntu 14.04 with November 2016 updates

Known Issues

There are no known issues in this 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 U.S. and other countries. Other company and product names may be trademarks of the respective companies with which they are associated.

Copyright

© 2020 NVIDIA Corporation. All rights reserved.

www.nvidia.com

