



NVCAFFE

RN-08464-001_v20.03 | April 2020

Release Notes



TABLE OF CONTENTS

Chapter 1. NVcaffe Overview.....	1
Chapter 2. Pulling a Container.....	2
Chapter 3. Running NVcaffe.....	3
3.1. Running An NVcaffe Container On A Cluster.....	4
Chapter 4. NVcaffe Release 20.03.....	5
Chapter 5. NVcaffe Release 20.02.....	8
Chapter 6. NVcaffe Release 20.01.....	11
Chapter 7. NVcaffe Release 19.12.....	14
Chapter 8. NVcaffe Release 19.11.....	16
Chapter 9. NVcaffe Release 19.10.....	18
Chapter 10. NVcaffe Release 19.09.....	20
Chapter 11. NVcaffe Release 19.08.....	22
Chapter 12. NVcaffe Release 19.07.....	24
Chapter 13. NVcaffe Release 19.06.....	26
Chapter 14. NVcaffe Release 19.05.....	28
Chapter 15. NVcaffe Release 19.04.....	30
Chapter 16. NVcaffe Release 19.03.....	32
Chapter 17. NVcaffe Release 19.02.....	34
Chapter 18. NVcaffe Release 19.01.....	36
Chapter 19. NVcaffe Release 18.12.....	38
Chapter 20. NVcaffe Release 18.11.....	40
Chapter 21. NVcaffe Release 18.10.....	42
Chapter 22. NVcaffe Release 18.09.....	44
Chapter 23. NVcaffe Release 18.08.....	46
Chapter 24. NVcaffe Release 18.07.....	48
Chapter 25. NVcaffe Release 18.06.....	50
Chapter 26. NVcaffe Release 18.05.....	52
Chapter 27. NVcaffe Release 18.04.....	53
Chapter 28. NVcaffe Release 18.03.....	55
Chapter 29. NVcaffe Release 18.02.....	57
Chapter 30. NVcaffe Release 18.01.....	59
Chapter 31. NVcaffe Release 17.12.....	61
Chapter 32. NVcaffe Release 17.11.....	62
Chapter 33. NVcaffe Release 17.10.....	63
Chapter 34. NVcaffe Release 17.09.....	64
Chapter 35. NVcaffe Release 17.07.....	66
Chapter 36. NVcaffe Release 17.06.....	67
Chapter 37. NVcaffe Release 17.05.....	68
Chapter 38. NVcaffe Release 17.04.....	69
Chapter 39. NVcaffe Release 17.03.....	70

Chapter 40. NVcaffe Release 17.02.....	71
Chapter 41. NVcaffe Release 17.01.....	72
Chapter 42. NVcaffe Release 16.12.....	73

Chapter 1.

NVCAFFE OVERVIEW

End of Life Announcement

Release 20.03 is the final release of the NVCAffe™ container.

The NVIDIA Deep Learning SDK accelerates widely-used deep learning frameworks such as NVCAffe.

Caffe™ was originally developed by the [Berkeley Vision and Learning Center \(BVLC\)](#) and by community contributors. It is a deep learning framework made with expression, speed, and modularity in mind.

NVCAffe is an NVIDIA-maintained fork of BVLC Caffe tuned for NVIDIA GPUs, particularly in multi-GPU configurations, accelerated by the NVIDIA Deep Learning SDK. It includes multi-precision support as well as other NVIDIA-enhanced features and offers performance specially tuned for NVIDIA GPU systems.

The NVCAffe framework is pre-built, tested, ready to run, and includes all of the necessary dependencies. The NVCAffe source code is also included for those who want to augment it with custom modifications or enhancements.

See `/workspace/README.md` inside the container for information on customizing your NVCAffe image. For more information about Caffe, including tutorials, documentation, and examples, see the [Caffe website](#).

NVCAffe typically utilizes the same input formats and configuration parameters as Caffe, therefore, community-authored materials and pre-trained models for Caffe can usually also be applied to NVCAffe.

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 be logged into the NGC container registry as explained in the [NGC Getting Started Guide](#).

The deep learning frameworks are stored in the following repository where you can find the NGC Docker containers.

`nvcr.io/nvidia`

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

Chapter 3.

RUNNING NVCAFFE

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, but can be specified using `NVIDIA_VISIBLE_DEVICES` environment variable). Starting in Docker 19.03, follow the steps as outlined below. For more information, refer to the nvidia-docker documentation [here](#).

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/caffe:20.03-py3
```

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:**

If you have Docker 19.03 or later, a typical command to launch the container is:

```
docker run --gpus all -it --rm -v local_dir:container_dir nvcr.io/nvidia/caffe:<xx.xx>-py3
```

If you have Docker 19.02 or earlier, a typical command to launch the container is:

```
nvidia-docker run -it --rm -v local_dir:container_dir nvcr.io/nvidia/caffe:<xx.xx>-py3
```

b) Non-interactive mode:

If you have Docker 19.03 or later, a typical command to launch the container is:

```
docker run --gpus all --rm -v local_dir:container_dir nvcr.io/nvidia/caffe:<xx.xx>-py3 caffe train ...
```

If you have Docker 19.02 or earlier, a typical command to launch the container is:

```
nvidia-docker run --rm -v local_dir:container_dir nvcr.io/nvidia/caffe:<xx.xx>-py3 caffe train ...
```

You might want to pull in data and model descriptions from locations outside the container for use by NVCaffe 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](#).



In order to share data between ranks, NVIDIA® Collective Communications Library™ (NCCL) may require shared system memory for IPC and pinned (page-locked) system memory resources. The operating system's limits on these resources may need to be increased accordingly. Refer to your system's documentation for details.

In particular, Docker containers default to limited shared and pinned memory resources. When using NCCL inside a container, it is recommended that you increase these resources by issuing:

```
--shm-size=1g --ulimit memlock=-1
```

in the command line to:

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

3.1. Running An NVCaffe Container On A Cluster

NVCaffe supports training on multiple nodes using [OpenMPI version 2.0](#) protocol, however, you cannot specify the number of threads per process because NVCaffe has its own thread manager (currently it runs one worker thread per GPU). For example:

```
dgx job submit --name jobname --volume <src>:<dst> --tasks 48
--clusterid <id> --gpu 8 --cpu 64 --mem 480 --image <tag> --nc "mpirun
-bind-to none -np 48 -pernode --tag-output caffe train --solver
solver.prototxt --gpu all >> /logs/caffe.log 2>&1"
```


Chapter 4.

NVCAFFE RELEASE 20.03

The NVIDIA container image of [Caffe](#), release 20.03, is available on [NGC](#).

Contents of the NVCAFFE container

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

The container also includes the following:

- ▶ [Ubuntu 18.04](#) including [Python 3.6](#)
- ▶ [NVIDIA CUDA 10.2.89](#) including [cuBLAS 10.2.2.89](#)
- ▶ [NVIDIA cuDNN 7.6.5](#)
- ▶ [NVIDIA NCCL 2.6.3](#) (optimized for NVLink™)
- ▶ [MLNX OFED](#)
- ▶ [OpenMPI 3.1.4](#)
- ▶ [Nsight Compute 2019.5.0](#)
- ▶ [Nsight Systems 2020.1.1](#)
- ▶ [TensorRT 7.0.0](#)
- ▶ Jupyter and JupyterLab:
 - ▶ [Jupyter Client 6.0.0](#)
 - ▶ [Jupyter Core 4.6.1](#)
 - ▶ [Jupyter Notebook 6.0.3](#)
 - ▶ [JupyterLab 1.1.1](#)
 - ▶ [JupyterLab Server 1.0.6](#)
 - ▶ [Jupyter-TensorBoard](#)

Driver Requirements

Release 20.03 is based on [NVIDIA CUDA 10.2.89](#), which requires [NVIDIA Driver](#) release 440.33.01. However, if you are running on Tesla (for example, T4 or any other Tesla board), you may use NVIDIA driver release 396, 384.111+, 410, 418.xx or 440.30. The CUDA driver's compatibility package only supports particular drivers. For a complete list of supported drivers, see the [CUDA Application Compatibility](#) topic. For more information, see [CUDA Compatibility and Upgrades](#).

GPU Requirements

Release 20.03 supports CUDA compute capability 6.0 and higher. This corresponds to GPUs in the Pascal, Volta, and Turing families. Specifically, for a list of GPUs that this compute capability corresponds to, see [CUDA GPUs](#). For additional support details, see [Deep Learning Frameworks Support Matrix](#).

Key Features and Enhancements

This NVCaffe release includes the following key features and enhancements.

- ▶ [NVCaffe 0.17.3](#) container image version 20.03 is based on [Caffe Deep Learning Framework by the BVLC](#) (<http://caffe.berkeleyvision.org/>).
- ▶ Ubuntu 18.04 with February 2020 updates

Announcements

- ▶ Deep learning framework containers 19.11 and later include experimental support for Singularity v3.0.

NVIDIA NVCaffe Container Versions

The following table shows what versions of Ubuntu, CUDA, NVCaffe, and TensorRT are supported in each of the NVIDIA containers for NVCaffe. For older container versions, refer to the [Frameworks Support Matrix](#).

Container Version	Ubuntu	CUDA Toolkit	NVCaffe	TensorRT
20.03	18.04	NVIDIA CUDA 10.2.89	0.17.3	TensorRT 7.0.0
20.02	16.04			
20.01				
19.12		NVIDIA CUDA 10.1.243		TensorRT 6.0.1
19.11				
19.10				

Container Version	Ubuntu	CUDA Toolkit	NVCaffe	TensorRT
19.09				
19.08				TensorRT 5.1.5

Known Issues

There are no known issues in this release.

Chapter 5.

NVCAFFE RELEASE 20.02

The NVIDIA container image of [Caffe](#), release 20.02, is available on [NGC](#).

Contents of the NVCAffe container

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

The container also includes the following:

- ▶ [Ubuntu 18.04](#) including [Python 3.6](#)
- ▶ [NVIDIA CUDA 10.2.89](#) including [cuBLAS 10.2.2.89](#)
- ▶ [NVIDIA cuDNN 7.6.5](#)
- ▶ [NVIDIA NCCL 2.5.6](#) (optimized for NVLink™)
- ▶ [MLNX OFED](#)
- ▶ [OpenMPI 3.1.4](#)
- ▶ [Nsight Compute 2019.5.0](#)
- ▶ [Nsight Systems 2020.1.1](#)
- ▶ [TensorRT 7.0.0](#)
- ▶ Jupyter and JupyterLab:
 - ▶ [Jupyter Client 5.3.4](#)
 - ▶ [Jupyter Core 4.6.1](#)
 - ▶ [Jupyter Notebook 6.0.3](#)
 - ▶ [JupyterLab 1.1.1](#)
 - ▶ [JupyterLab Server 1.0.6](#)
 - ▶ [Jupyter-TensorBoard](#)

Driver Requirements

Release 20.02 is based on [NVIDIA CUDA 10.2.89](#), which requires [NVIDIA Driver](#) release 440.33.01. However, if you are running on Tesla (for example, T4 or any other Tesla board), you may use NVIDIA driver release 396, 384.111+, 410, 418.xx or 440.30. The CUDA driver's compatibility package only supports particular drivers. For a complete list of supported drivers, see the [CUDA Application Compatibility](#) topic. For more information, see [CUDA Compatibility and Upgrades](#).

GPU Requirements

Release 20.02 supports CUDA compute capability 6.0 and higher. This corresponds to GPUs in the Pascal, Volta, and Turing families. Specifically, for a list of GPUs that this compute capability corresponds to, see [CUDA GPUs](#). For additional support details, see [Deep Learning Frameworks Support Matrix](#).

Key Features and Enhancements

This NVCaffe release includes the following key features and enhancements.

- ▶ [NVCaffe 0.17.3](#) container image version 20.02 is based on [Caffe Deep Learning Framework](#) by the BVLC (<http://caffe.berkeleyvision.org/>).
- ▶ The latest version of [Nsight Systems 2020.1.1](#)
- ▶ The latest version of [Jupyter Notebook 6.0.3](#)
- ▶ Ubuntu 18.04 with January 2020 updates

Announcements

- ▶ Deep learning framework containers 19.11 and later include experimental support for Singularity v3.0.

NVIDIA NVCaffe Container Versions

The following table shows what versions of Ubuntu, CUDA, NVCaffe, and TensorRT are supported in each of the NVIDIA containers for NVCaffe. For older container versions, refer to the [Frameworks Support Matrix](#).

Container Version	Ubuntu	CUDA Toolkit	NVCaffe	TensorRT
20.02	18.04	NVIDIA CUDA 10.2.89	0.17.3	TensorRT 7.0.0
20.01	16.04			
19.12				TensorRT 6.0.1
19.11				

Container Version	Ubuntu	CUDA Toolkit	NVCaffe	TensorRT
19.10		NVIDIA CUDA 10.1.243		
19.09				
19.08				TensorRT 5.1.5

Known Issues

There are no known issues in this release.

Chapter 6.

NVCAFFE RELEASE 20.01

The NVIDIA container image of [Caffe](#), release 20.01, is available on [NGC](#).

Contents of the NVCAffe container

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

The container also includes the following:

- ▶ [Ubuntu 18.04](#) including [Python 3.6](#)
- ▶ [NVIDIA CUDA 10.2.89](#) including [cuBLAS 10.2.2.89](#)
- ▶ [NVIDIA cuDNN 7.6.5](#)
- ▶ [NVIDIA NCCL 2.5.6](#) (optimized for NVLink™)
- ▶ [MLNX OFED](#)
- ▶ [OpenMPI 3.1.4](#)
- ▶ [Nsight Compute 2019.5.0](#)
- ▶ [Nsight Systems 2019.6.1](#)
- ▶ [TensorRT 7.0.0](#)
- ▶ Jupyter and JupyterLab:
 - ▶ [Jupyter Client 5.3.4](#)
 - ▶ [Jupyter Core 4.6.1](#)
 - ▶ [Jupyter Notebook 6.0.1](#)
 - ▶ [JupyterLab 1.1.1](#)
 - ▶ [JupyterLab Server 1.0.6](#)
 - ▶ [Jupyter-TensorBoard](#)

Driver Requirements

Release 20.01 is based on [NVIDIA CUDA 10.2.89](#), which requires [NVIDIA Driver](#) release 440.33.01. However, if you are running on Tesla (for example, T4 or any other Tesla board), you may use NVIDIA driver release 396, 384.111+, 410, 418.xx or 440.30. The CUDA driver's compatibility package only supports particular drivers. For a complete list of supported drivers, see the [CUDA Application Compatibility](#) topic. For more information, see [CUDA Compatibility and Upgrades](#).

GPU Requirements

Release 20.01 supports CUDA compute capability 6.0 and higher. This corresponds to GPUs in the Pascal, Volta, and Turing families. Specifically, for a list of GPUs that this compute capability corresponds to, see [CUDA GPUs](#). For additional support details, see [Deep Learning Frameworks Support Matrix](#).

Key Features and Enhancements

This NVCaffe release includes the following key features and enhancements.

- ▶ [NVCaffe 0.17.3](#) container image version 20.01 is based on [Caffe Deep Learning Framework by the BVLC](#) (<http://caffe.berkeleyvision.org/>).
- ▶ Latest version of [TensorRT 7.0.0](#)
- ▶ Ubuntu 18.04 with December 2019 updates

Announcements

- ▶ Deep learning framework containers 19.11 and later include experimental support for Singularity v3.0.

NVIDIA NVCaffe Container Versions

The following table shows what versions of Ubuntu, CUDA, NVCaffe, and TensorRT are supported in each of the NVIDIA containers for NVCaffe. For older container versions, refer to the [Frameworks Support Matrix](#).

Container Version	Ubuntu	CUDA Toolkit	NVCaffe	TensorRT
20.01	18.04	NVIDIA CUDA 10.2.89	0.17.3	TensorRT 7.0.0
19.12	16.04			TensorRT 6.0.1
19.11				
19.10		NVIDIA CUDA 10.1.243		
19.09				
19.08				TensorRT 5.1.5

Known Issues

There are no known issues in this release.

Chapter 7.

NVCAFFE RELEASE 19.12

The NVIDIA container image of [Caffe](#), release 19.12, is available on [NGC](#).

Contents of the NVCAffe container

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

The container also includes the following:

- ▶ [Ubuntu 18.04](#) including [Python 3.6](#)
- ▶ [NVIDIA CUDA 10.2.89](#) including [cuBLAS 10.2.2.89](#)
- ▶ [NVIDIA cuDNN 7.6.5](#)
- ▶ [NVIDIA NCCL 2.5.6](#) (optimized for NVLink™)
- ▶ [MLNX OFED](#)
- ▶ [OpenMPI 3.1.4](#)
- ▶ [Nsight Compute 2019.5.0](#)
- ▶ [Nsight Systems 2019.6.1](#)
- ▶ [TensorRT 6.0.1](#)
- ▶ Jupyter and JupyterLab:
 - ▶ [Jupyter Client 5.3.4](#)
 - ▶ [Jupyter Core 4.6.1](#)
 - ▶ [Jupyter Notebook 6.0.1](#)
 - ▶ [JupyterLab 1.1.1](#)
 - ▶ [JupyterLab Server 1.0.6](#)
 - ▶ [Jupyter-TensorBoard](#)

Driver Requirements

Release 19.12 is based on [NVIDIA CUDA 10.2.89](#), which requires [NVIDIA Driver](#) release 440.30. However, if you are running on Tesla (for example, T4 or any other Tesla board), you may use NVIDIA driver release 396, 384.111+, 410, 418.xx or 440.30. The CUDA driver's compatibility package only supports particular drivers. For a complete list of supported drivers, see the [CUDA Application Compatibility](#) topic. For more information, see [CUDA Compatibility and Upgrades](#).

GPU Requirements

Release 19.12 supports CUDA compute capability 6.0 and higher. This corresponds to GPUs in the Pascal, Volta, and Turing families. Specifically, for a list of GPUs that this compute capability corresponds to, see [CUDA GPUs](#). For additional support details, see [Deep Learning Frameworks Support Matrix](#).

Key Features and Enhancements

This NVCaffe release includes the following key features and enhancements.

- ▶ [NVCaffe 0.17.3](#) container image version 19.12 is based on [Caffe Deep Learning Framework by the BVLC](#) (<http://caffe.berkeleyvision.org/>).
- ▶ Latest version of [Nsight Systems 2019.6.1](#)
- ▶ Ubuntu 18.04 with November 2019 updates

Announcements

- ▶ Deep learning framework containers 19.11 and later include experimental support for Singularity v3.0.

Known Issues

There are no known issues in this release.

Chapter 8.

NVCAFFE RELEASE 19.11

The NVIDIA container image of [Caffe](#), release 19.11, is available on [NGC](#).

Contents of the NVCAFFE container

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

The container also includes the following:

- ▶ [Ubuntu 18.04](#) including [Python 3.6](#)
- ▶ [NVIDIA CUDA 10.2.89](#) including [cuBLAS 10.2.2.89](#)
- ▶ [NVIDIA cuDNN 7.6.5](#)
- ▶ [NVIDIA NCCL 2.5.6](#) (optimized for NVLink™)
- ▶ [MLNX OFED](#)
- ▶ [OpenMPI 3.1.4](#)
- ▶ [Nsight Compute 2019.5.0](#)
- ▶ [Nsight Systems 2019.5.2](#)
- ▶ [TensorRT 6.0.1](#)
- ▶ Jupyter and JupyterLab:
 - ▶ [Jupyter Client 5.3.4](#)
 - ▶ [Jupyter Core 4.6.1](#)
 - ▶ [Jupyter Notebook 6.0.1](#)
 - ▶ [JupyterLab 1.1.1](#)
 - ▶ [JupyterLab Server 1.0.6](#)
 - ▶ [Jupyter-TensorBoard](#)

Driver Requirements

Release 19.11 is based on [NVIDIA CUDA 10.2.89](#), which requires [NVIDIA Driver](#) release 440.30. However, if you are running on Tesla (for example, T4 or any other Tesla board), you may use NVIDIA driver release 396, 384.111+, 410 or 418.xx. The CUDA driver's compatibility package only supports particular drivers. For a complete list of supported drivers, see the [CUDA Application Compatibility](#) topic. For more information, see [CUDA Compatibility and Upgrades](#).

GPU Requirements

Release 19.11 supports CUDA compute capability 6.0 and higher. This corresponds to GPUs in the Pascal, Volta, and Turing families. Specifically, for a list of GPUs that this compute capability corresponds to, see [CUDA GPUs](#). For additional support details, see [Deep Learning Frameworks Support Matrix](#).

Key Features and Enhancements

This NVCaffe release includes the following key features and enhancements.

- ▶ [NVCaffe 0.17.3](#) container image version 19.11 is based on [Caffe Deep Learning Framework by the BVLC](#) (<http://caffe.berkeleyvision.org/>).
- ▶ Latest version of [Python 3.6](#)
- ▶ Latest version of [NVIDIA CUDA 10.2.89](#) including [cuBLAS 10.2.2.89](#)
- ▶ Latest version of [NVIDIA cuDNN 7.6.5](#)
- ▶ Latest version of [NVIDIA NCCL 2.5.6](#)
- ▶ Latest version of [Nsight Compute 2019.5.0](#)
- ▶ Latest version of [Nsight Systems 2019.5.2](#)
- ▶ Latest versions of [Jupyter Client 5.3.4](#) and [Jupyter Core 4.6.1](#)
- ▶ Ubuntu 18.04 with October 2019 updates

Announcements

- ▶ Deep learning framework containers 19.11 and later include experimental support for Singularity v3.0.

Known Issues

There are no known issues in this release.

Chapter 9.

NVCAFFE RELEASE 19.10

The NVIDIA container image of [Caffe](#), release 19.10, is available on [NGC](#).

Contents of the NVCAFFE container

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

The container also includes the following:

- ▶ [Ubuntu 18.04](#) including [Python 2.7](#)
- ▶ [NVIDIA CUDA 10.1.243](#) including [cuBLAS 10.2.1.243](#)
- ▶ [NVIDIA cuDNN 7.6.4](#)
- ▶ [NVIDIA NCCL 2.4.8](#) (optimized for NVLink™)
- ▶ [MLNX OFED](#)
- ▶ [OpenMPI 3.1.4](#)
- ▶ [Nsight Compute 2019.4.0](#)
- ▶ [Nsight Systems 2019.5.1](#)
- ▶ [TensorRT 6.0.1](#)
- ▶ Jupyter and JupyterLab:
 - ▶ [Jupyter Client 5.3.3](#)
 - ▶ [Jupyter Core 4.5.0](#)
 - ▶ [Jupyter Notebook 6.0.1](#)
 - ▶ [JupyterLab 1.1.1](#)
 - ▶ [JupyterLab Server 1.0.6](#)
 - ▶ [Jupyter-TensorBoard](#)

Driver Requirements

Release 19.10 is based on [NVIDIA CUDA 10.1.243](#), which requires [NVIDIA Driver](#) release 418.xx. However, if you are running on Tesla (for example, T4 or any other Tesla board), you may use NVIDIA driver release 396, 384.111+ or 410. The CUDA driver's compatibility package only supports particular drivers. For a complete list of supported drivers, see the [CUDA Application Compatibility](#) topic. For more information, see [CUDA Compatibility and Upgrades](#).

GPU Requirements

Release 19.10 supports CUDA compute capability 6.0 and higher. This corresponds to GPUs in the Pascal, Volta, and Turing families. Specifically, for a list of GPUs that this compute capability corresponds to, see [CUDA GPUs](#). For additional support details, see [Deep Learning Frameworks Support Matrix](#).

Key Features and Enhancements

This NVCaffe release includes the following key features and enhancements.

- ▶ [NVCaffe 0.17.3](#) container image version 19.10 is based on [Caffe Deep Learning Framework by the BVLC](#) (<http://caffe.berkeleyvision.org/>).
- ▶ Latest version of [NVIDIA cuDNN 7.6.4](#)
- ▶ Latest version of [Nsight Systems 2019.5.1](#)
- ▶ Latest version of [Jupyter Client 5.3.3](#)
- ▶ Ubuntu 18.04 with September 2019 updates

Known Issues

There are no known issues in this release.

Chapter 10.

NVCAFFE RELEASE 19.09

The NVIDIA container image of [Caffe](#), release 19.09, is available on [NGC](#).

Contents of the NVCAFFE container

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

The container also includes the following:

- ▶ [Ubuntu 18.04](#) including [Python 2.7](#)
- ▶ [NVIDIA CUDA 10.1.243](#) including [cuBLAS 10.2.1.243](#)
- ▶ [NVIDIA cuDNN 7.6.3](#)
- ▶ [NVIDIA NCCL 2.4.8](#) (optimized for NVLink™)
- ▶ [MLNX OFED](#)
- ▶ [OpenMPI 3.1.4](#)
- ▶ [Nsight Compute 2019.4.0](#)
- ▶ [Nsight Systems 2019.4.2](#)
- ▶ [TensorRT 6.0.1](#)
- ▶ Jupyter and JupyterLab:
 - ▶ [Jupyter Client 5.3.1](#)
 - ▶ [Jupyter Core 4.5.0](#)
 - ▶ [Jupyter Notebook 6.0.1](#)
 - ▶ [JupyterLab 1.1.1](#)
 - ▶ [JupyterLab Server 1.0.6](#)
 - ▶ [Jupyter-TensorBoard](#)

Driver Requirements

Release 19.09 is based on [NVIDIA CUDA 10.1.243](#), which requires [NVIDIA Driver](#) release 418.xx. However, if you are running on Tesla (for example, T4 or any other Tesla board), you may use NVIDIA driver release 396, 384.111+ or 410. The CUDA driver's compatibility package only supports particular drivers. For a complete list of supported drivers, see the [CUDA Application Compatibility](#) topic. For more information, see [CUDA Compatibility and Upgrades](#).

GPU Requirements

Release 19.09 supports CUDA compute capability 6.0 and higher. This corresponds to GPUs in the Pascal, Volta, and Turing families. Specifically, for a list of GPUs that this compute capability corresponds to, see [CUDA GPUs](#). For additional support details, see [Deep Learning Frameworks Support Matrix](#).

Key Features and Enhancements

This NVCaffe release includes the following key features and enhancements.

- ▶ [NVCaffe 0.17.3](#) container image version 19.09 is based on [Caffe Deep Learning Framework](#) by the BVLC (<http://caffe.berkeleyvision.org/>).
- ▶ Latest version of [NVIDIA cuDNN 7.6.3](#)
- ▶ Latest version of [TensorRT 6.0.1](#)
- ▶ Latest versions of [Nsight Compute 2019.4.0](#) and [Nsight Systems 2019.4.2](#)
- ▶ Latest versions of [Jupyter Notebook 6.0.1](#), [JupyterLab 1.1.1](#), and [JupyterLab Server 1.0.6](#).
- ▶ Ubuntu 18.04 with August 2019 updates

Known Issues

There are no known issues in this release.

Chapter 11.

NVCAFFE RELEASE 19.08

The NVIDIA container image of [Caffe](#), release 19.08, is available on [NGC](#).

Contents of the NVCAFFE container

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

The container also includes the following:

- ▶ [Ubuntu 18.04](#) including [Python 2.7](#)
- ▶ [NVIDIA CUDA 10.1.243](#) including [cuBLAS 10.2.1.243](#)
- ▶ [NVIDIA cuDNN 7.6.2](#)
- ▶ [NVIDIA NCCL 2.4.8](#) (optimized for NVLink™)
- ▶ [MLNX OFED +4.0](#)
- ▶ [OpenMPI 3.1.4](#)
- ▶ [Nsight Compute 10.1.168](#)
- ▶ [Nsight Systems 2019.3.7.9](#)
- ▶ [TensorRT 5.1.5](#)
- ▶ Jupyter and JupyterLab:
 - ▶ [Jupyter Client 5.3.1](#)
 - ▶ [Jupyter Core 4.5.0](#)
 - ▶ [Jupyter Notebook 6.0.0](#)
 - ▶ [JupyterLab 1.0.5](#)
 - ▶ [JupyterLab Server 1.0.0](#)
 - ▶ [Jupyter-TensorBoard](#)

Driver Requirements

Release 19.08 is based on [NVIDIA CUDA 10.1.243](#), which requires [NVIDIA Driver](#) release 418.87. However, if you are running on Tesla (Tesla V100, Tesla P4, Tesla P40, or Tesla P100), you may use NVIDIA driver release 384.111+ or 410. The CUDA driver's compatibility package only supports particular drivers. For a complete list of supported drivers, see the [CUDA Application Compatibility](#) topic. For more information, see [CUDA Compatibility and Upgrades](#).

GPU Requirements

Release 19.08 supports CUDA compute capability 6.0 and higher. This corresponds to GPUs in the Pascal, Volta, and Turing families. Specifically, for a list of GPUs that this compute capability corresponds to, see [CUDA GPUs](#). For additional support details, see [Deep Learning Frameworks Support Matrix](#).

Key Features and Enhancements

This NVCaffe release includes the following key features and enhancements.

- ▶ [NVCaffe 0.17.3](#) container image version 19.08 is based on [Caffe Deep Learning Framework](#) by the BVLC (<http://caffe.berkeleyvision.org/>).
- ▶ Latest version of [NVIDIA CUDA 10.1.243](#) including [cuBLAS 10.2.1.243](#)
- ▶ Latest version of [NVIDIA cuDNN 7.6.2](#)
- ▶ Latest version of [NVIDIA NCCL 2.4.8](#)
- ▶ Latest version of [OpenMPI 3.1.4](#)
- ▶ Latest version of [Nsight Systems 2019.3.7.9](#)
- ▶ Latest version of [MLNX OFED +4.0](#)
- ▶ Latest versions of [Jupyter Notebook 6.0.0](#) and [JupyterLab 1.0.5](#)
- ▶ Ubuntu 18.04 with July 2019 updates

Known Issues

There are no known issues in this release.

Chapter 12.

NVCAFFE RELEASE 19.07

The NVIDIA container image of [Caffe](#), release 19.07, is available on [NGC](#).

Contents of the NVCAffe container

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

The container also includes the following:

- ▶ [Ubuntu 18.04](#) including [Python 2.7](#)
- ▶ [NVIDIA CUDA 10.1.168](#) including [cuBLAS 10.2.0.168](#)
- ▶ [NVIDIA cuDNN 7.6.1](#)
- ▶ [NVIDIA NCCL 2.4.7](#) (optimized for NVLink™)
- ▶ [MLNX OFED +3.4](#)
- ▶ [OpenMPI 3.1.3](#)
- ▶ [TensorRT 5.1.5](#)
- ▶ Jupyter and JupyterLab:
 - ▶ [Jupyter Client 5.3.1](#)
 - ▶ [Jupyter Core 4.5.0](#)
 - ▶ [Jupyter Notebook 5.7.8](#)
 - ▶ [JupyterLab 1.0.2](#)
 - ▶ [JupyterLab Server 1.0.0](#)
 - ▶ [Jupyter-TensorBoard](#)

Driver Requirements

Release 19.07 is based on [NVIDIA CUDA 10.1.168](#), which requires [NVIDIA Driver](#) release 418.67. However, if you are running on Tesla (Tesla V100, Tesla P4, Tesla P40, or Tesla P100), you may use NVIDIA driver release 384.111+ or 410. The CUDA driver's

compatibility package only supports particular drivers. For a complete list of supported drivers, see the [CUDA Application Compatibility](#) topic. For more information, see [CUDA Compatibility and Upgrades](#).

GPU Requirements

Release 19.07 supports CUDA compute capability 6.0 and higher. This corresponds to GPUs in the Pascal, Volta, and Turing families. Specifically, for a list of GPUs that this compute capability corresponds to, see [CUDA GPUs](#). For additional support details, see [Deep Learning Frameworks Support Matrix](#).

Key Features and Enhancements

This NVCaffe release includes the following key features and enhancements.

- ▶ [NVCaffe 0.17.3](#) container image version 19.07 is based on [Caffe Deep Learning Framework by the BVLC \(http://caffe.berkeleyvision.org/\)](#).
- ▶ Fixed a bug that prevented the usage of batch size = 1 with pictures of different sizes.
- ▶ Latest version of [NVIDIA cuDNN 7.6.1](#)
- ▶ Latest version of [MLNX OFED +3.4](#)
- ▶ Latest versions of [Jupyter Client 5.3.1](#), [Jupyter Core 4.5.0](#), [JupyterLab 1.0.2](#) and [JupyterLab Server 1.0.0](#), including [Jupyter-TensorBoard](#) integration.
- ▶ Latest version of [Ubuntu 18.04](#)

Known Issues

There are no known issues in this release.

Chapter 13.

NVCAFFE RELEASE 19.06

The NVIDIA container image of [Caffe](#), release 19.06, is available on [NGC](#).

Contents of the NVCAffe container

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

The container also includes the following:

- ▶ [Ubuntu 16.04](#) including [Python 2.7](#)
- ▶ [NVIDIA CUDA 10.1.168](#) including [cuBLAS 10.2.0.168](#)
- ▶ [NVIDIA cuDNN 7.6.0](#)
- ▶ [NVIDIA NCCL 2.4.7](#) (optimized for NVLink™)
- ▶ [OpenMPI 3.1.3](#)
- ▶ [TensorRT 5.1.5](#)
- ▶ Jupyter and JupyterLab:
 - ▶ [Jupyter Client 5.2.4](#)
 - ▶ [Jupyter Core 4.4.0](#)
 - ▶ [Jupyter Notebook 5.7.8](#)
 - ▶ [JupyterLab 0.35.4](#)
 - ▶ [JupyterLab Server 0.2.0](#)

Driver Requirements

Release 19.06 is based on [NVIDIA CUDA 10.1.168](#), which requires [NVIDIA Driver](#) release 418.xx. However, if you are running on Tesla (Tesla V100, Tesla P4, Tesla P40, or Tesla P100), you may use NVIDIA driver release 384.111+ or 410. The CUDA driver's compatibility package only supports particular drivers. For a complete list of supported drivers, see the [CUDA Application Compatibility](#) topic. For more information, see [CUDA Compatibility and Upgrades](#).

GPU Requirements

Release 19.06 supports CUDA compute capability 6.0 and higher. This corresponds to GPUs in the Pascal, Volta, and Turing families. Specifically, for a list of GPUs that this compute capability corresponds to, see [CUDA GPUs](#). For additional support details, see [Deep Learning Frameworks Support Matrix](#).

Key Features and Enhancements

This NVCaffe release includes the following key features and enhancements.

- ▶ [NVCaffe 0.17.3](#) container image version 19.06 is based on [Caffe Deep Learning Framework by the BVLC](#) (<http://caffe.berkeleyvision.org/>).
- ▶ Latest version of [NVIDIA CUDA 10.1.168](#) including [cuBLAS 10.2.0.168](#)
- ▶ Latest version of [NVIDIA NCCL 2.4.7](#)
- ▶ Latest version of [JupyterLab 0.35.6](#)
- ▶ Ubuntu 16.04 with May 2019 updates (see Announcements)

Announcements

In the next release, we will no longer support [Ubuntu 16.04](#). Release 19.07 will instead support [Ubuntu 18.04](#).

Known Issues

There are no known issues in this release.

Chapter 14.

NVCAFFE RELEASE 19.05

The NVIDIA container image of [Caffe](#), release 19.05, is available on [NGC](#).

Contents of the NVCAffe container

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

The container also includes the following:

- ▶ [Ubuntu 16.04](#) including [Python 2.7](#)
- ▶ [NVIDIA CUDA 10.1 Update 1](#) including [cuBLAS 10.1 Update 1](#)
- ▶ [NVIDIA cuDNN 7.6.0](#)
- ▶ [NVIDIA NCCL 2.4.6](#) (optimized for NVLink™)
- ▶ [OpenMPI 3.1.3](#)
- ▶ [TensorRT 5.1.5](#)
- ▶ Jupyter and JupyterLab:
 - ▶ Jupyter Client 5.2.4
 - ▶ Jupyter Core 4.4.0
 - ▶ JupyterLab 0.35.4
 - ▶ JupyterLab Server 0.2.0

Driver Requirements

Release 19.05 is based on CUDA 10.1 Update 1, which requires [NVIDIA Driver](#) release 418.xx. However, if you are running on Tesla (Tesla V100, Tesla P4, Tesla P40, or Tesla P100), you may use NVIDIA driver release 384.111+ or 410. The CUDA driver's compatibility package only supports particular drivers. For a complete list of supported drivers, see the [CUDA Application Compatibility](#) topic. For more information, see [CUDA Compatibility and Upgrades](#).

GPU Requirements

Release 19.05 supports CUDA compute capability 6.0 and higher. This corresponds to GPUs in the Pascal, Volta, and Turing families. Specifically, for a list of GPUs that this compute capability corresponds to, see [CUDA GPUs](#). For additional support details, see [Deep Learning Frameworks Support Matrix](#).

Key Features and Enhancements

This NVCaffe release includes the following key features and enhancements.

- ▶ [NVCaffe 0.17.3](#) container image version 19.05 is based on [Caffe Deep Learning Framework by the BVLC](#) (<http://caffe.berkeleyvision.org/>).
- ▶ Latest version of [NVIDIA CUDA 10.1 Update 1](#) including [cuBLAS 10.1 Update 1](#)
- ▶ Latest version of [NVIDIA cuDNN 7.6.0](#)
- ▶ Latest version of [TensorRT 5.1.5](#)
- ▶ Latest version of [DALI 0.9.1 Beta](#)
- ▶ Ubuntu 16.04 with April 2019 updates

Known Issues

There are no known issues in this release.

Chapter 15.

NVCAFFE RELEASE 19.04

The NVIDIA container image of [Caffe](#), release 19.04, is available on [NGC](#).

Contents of the NVCAFFE container

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

The container also includes the following:

- ▶ [Ubuntu 16.04](#) including [Python 2.7](#)
- ▶ [NVIDIA CUDA 10.1.105](#) including [cuBLAS 10.1.0.105](#)
- ▶ [NVIDIA cuDNN 7.5.0](#)
- ▶ [NVIDIA NCCL 2.4.6](#) (optimized for NVLink™)
- ▶ [OpenMPI 3.1.3](#)
- ▶ [TensorRT 5.1.2](#)
- ▶ Jupyter and JupyterLab:
 - ▶ Jupyter Client 5.2.4
 - ▶ Jupyter Core 4.4.0
 - ▶ JupyterLab 0.35.4
 - ▶ JupyterLab Server 0.2.0

Driver Requirements

Release 19.04 is based on CUDA 10.1, which requires [NVIDIA Driver](#) release 418.xx.x+. However, if you are running on Tesla (Tesla V100, Tesla P4, Tesla P40, or Tesla P100), you may use NVIDIA driver release 384.111+ or 410. The CUDA driver's compatibility package only supports particular drivers. For a complete list of supported drivers, see the [CUDA Application Compatibility](#) topic. For more information, see [CUDA Compatibility and Upgrades](#).

GPU Requirements

Release 19.04 supports CUDA compute capability 6.0 and higher. This corresponds to GPUs in the Pascal, Volta, and Turing families. Specifically, for a list of GPUs that this compute capability corresponds to, see [CUDA GPUs](#). For additional support details, see [Deep Learning Frameworks Support Matrix](#).

Key Features and Enhancements

This NVCaffe release includes the following key features and enhancements.

- ▶ [NVCaffe 0.17.3](#) container image version 19.04 is based on [Caffe Deep Learning Framework by the BVLC](#) (<http://caffe.berkeleyvision.org/>).
- ▶ Latest version of [NVIDIA NCCL 2.4.6](#)
- ▶ Latest version of [cuBLAS 10.1.0.105](#)
- ▶ Ubuntu 16.04 with March 2019 updates

Known Issues

There are no known issues in this release.

Chapter 16.

NVCAFFE RELEASE 19.03

The NVIDIA container image of [Caffe](#), release 19.03, is available on [NGC](#).

Contents of the NVCAffe container

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

The container also includes the following:

- ▶ [Ubuntu 16.04](#) including [Python 2.7](#)
- ▶ [NVIDIA CUDA 10.1.105](#) including [cuBLAS 10.1.105](#)
- ▶ [NVIDIA cuDNN 7.5.0](#)
- ▶ [NVIDIA NCCL 2.4.3](#) (optimized for NVLink™)
- ▶ [OpenMPI 3.1.3](#)
- ▶ [TensorRT 5.1.2](#)
- ▶ Jupyter and JupyterLab:
 - ▶ Jupyter Client 5.2.4
 - ▶ Jupyter Core 4.4.0
 - ▶ JupyterLab 0.35.4
 - ▶ JupyterLab Server 0.2.0

Driver Requirements

Release 19.03 is based on CUDA 10.1, which requires [NVIDIA Driver](#) release 418.xx+. However, if you are running on Tesla (Tesla V100, Tesla P4, Tesla P40, or Tesla P100), you may use NVIDIA driver release 384.111+ or 410. The CUDA driver's compatibility package only supports particular drivers. For a complete list of supported drivers, see the [CUDA Application Compatibility](#) topic. For more information, see [CUDA Compatibility and Upgrades](#).

GPU Requirements

Release 19.03 supports CUDA compute capability 6.0 and higher. This corresponds to GPUs in the Pascal, Volta, and Turing families. Specifically, for a list of GPUs that this compute capability corresponds to, see [CUDA GPUs](#). For additional support details, see [Deep Learning Frameworks Support Matrix](#).

Key Features and Enhancements

This NVCaffe release includes the following key features and enhancements.

- ▶ [NVCaffe 0.17.3](#) container image version 19.03 is based on [Caffe Deep Learning Framework by the BVLC](#) (<http://caffe.berkeleyvision.org/>).
- ▶ Latest version of [NVIDIA CUDA 10.1.105](#) including [cuBLAS 10.1.105](#)
- ▶ Latest version of [NVIDIA cuDNN 7.5.0](#)
- ▶ Latest version of [NVIDIA NCCL 2.4.3](#)
- ▶ Latest version of [TensorRT 5.1.2](#)
- ▶ Ubuntu 16.04 with February 2019 updates

Known Issues

- ▶ If using or upgrading to a 3-part-version driver, for example, a driver that takes the format of **xxx.yy.zz**, you will receive a **Failed to detect NVIDIA driver version.** message. This is due to a known bug in the entry point script's parsing of 3-part driver versions. This message is non-fatal and can be ignored. This will be fixed in the 19.04 release.

Chapter 17.

NVCAFFE RELEASE 19.02

The NVIDIA container image of [Caffe](#), release 19.02 is available.

Contents of NVCAffe

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

The container also includes the following:

- ▶ [Ubuntu 16.04](#) including [Python 2.7](#)
- ▶ [NVIDIA CUDA 10.0.130](#) including [CUDA[®] Basic Linear Algebra Subroutines library[™] \(cuBLAS\) 10.0.130](#)
- ▶ [NVIDIA CUDA[®] Deep Neural Network library[™] \(cuDNN\) 7.4.2](#)
- ▶ [NVIDIA Collective Communications Library \(NCCL\) 2.3.7](#) (optimized for [NVLink[™]](#))
- ▶ [OpenMPI 3.1.3](#)
- ▶ [TensorRT 5.0.2](#)
- ▶ Jupyter and JupyterLab:
 - ▶ [Jupyter Client 5.2.4](#)
 - ▶ [Jupyter Core 4.4.0](#)
 - ▶ [JupyterLab 0.35.4](#)
 - ▶ [JupyterLab Server 0.2.0](#)

Driver Requirements

Release 19.02 is based on CUDA 10, which requires [NVIDIA Driver](#) release 410.xx. However, if you are running on Tesla (Tesla V100, Tesla P4, Tesla P40, or Tesla P100), you may use NVIDIA driver release 384. For more information, see [CUDA Compatibility and Upgrades](#).

GPU Requirements

Release 19.02 supports CUDA compute capability 6.0 and higher. This corresponds to GPUs in the Pascal, Volta, and Turing families. Specifically, for a list of GPUs that this compute capability corresponds to, see [CUDA GPUs](#). For additional support details, see [Deep Learning Frameworks Support Matrix](#).

Key Features and Enhancements

This NVCaffe release includes the following key features and enhancements.

- ▶ NVCaffe 0.17.2 container image version 19.02 is based on [Caffe Deep Learning Framework by the BVLC \(http://caffe.berkeleyvision.org/\)](#).
- ▶ Added Jupyter and JupyterLab software in our packaged container.
- ▶ Latest version of [jupyter client 5.2.4](#)
- ▶ Latest version of [jupyter core 4.4.0](#)
- ▶ Ubuntu 16.04 with January 2019 updates

Known Issues

- ▶ There is a known performance regression with Caffe 0.17.2 for some networks when run with small batch sizes. As a workaround, increase the batch size.
- ▶ If using or upgrading to a 3-part-version driver, for example, a driver that takes the format of **xxx.yy.zz**, you will receive a **Failed to detect NVIDIA driver version.** message. This is due to a known bug in the entry point script's parsing of 3-part driver versions. This message is non-fatal and can be ignored. This will be fixed in the 19.04 release.

Chapter 18.

NVCAFFE RELEASE 19.01

The NVIDIA container image of [Caffe](#), release 19.01 is available.

Contents of NVCAffe

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

The container also includes the following:

- ▶ [Ubuntu 16.04](#) including [Python 2.7](#)
- ▶ [NVIDIA CUDA 10.0.130](#) including [CUDA[®] Basic Linear Algebra Subroutines library[™] \(cuBLAS\) 10.0.130](#)
- ▶ [NVIDIA CUDA[®] Deep Neural Network library[™] \(cuDNN\) 7.4.2](#)
- ▶ [NCCL 2.3.7](#) (optimized for [NVLink[™]](#))
- ▶ [OpenMPI 3.1.3](#)
- ▶ [TensorRT 5.0.2](#)

Driver Requirements

Release 19.01 is based on CUDA 10, which requires [NVIDIA Driver](#) release 410.xx. However, if you are running on Tesla (Tesla V100, Tesla P4, Tesla P40, or Tesla P100), you may use NVIDIA driver release 384. For more information, see [CUDA Compatibility and Upgrades](#).

GPU Requirements

Release 19.01 supports CUDA compute capability 6.0 and higher. This corresponds to GPUs in the Pascal, Volta, and Turing families. Specifically, for a list of GPUs that this compute capability corresponds to, see [CUDA GPUs](#). For additional support details, see [Deep Learning Frameworks Support Matrix](#).

Key Features and Enhancements

This NVCaffe release includes the following key features and enhancements.

- ▶ NVCaffe 0.17.2 container image version 19.01 is based on [Caffe Deep Learning Framework by the BVLC](http://caffe.berkeleyvision.org/) (<http://caffe.berkeleyvision.org/>).
- ▶ Latest version of [NVIDIA cuDNN 7.4.2](#)
- ▶ Ubuntu 16.04 with December 2018 updates

Known Issues

- ▶ If using or upgrading to a 3-part-version driver, for example, a driver that takes the format of **xxx.yy.zz**, you will receive a **Failed to detect NVIDIA driver version.** message. This is due to a known bug in the entry point script's parsing of 3-part driver versions. This message is non-fatal and can be ignored. This will be fixed in the 19.04 release.

Chapter 19.

NVCAFFE RELEASE 18.12

The NVIDIA container image of [Caffe](#), release 18.12 is available.

Contents of NVCaffe

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

The container also includes the following:

- ▶ [Ubuntu 16.04](#) including [Python 2.7](#)
- ▶ [NVIDIA CUDA 10.0.130](#) including [CUDA[®] Basic Linear Algebra Subroutines library[™] \(cuBLAS\) 10.0.130](#)
- ▶ [NVIDIA CUDA[®] Deep Neural Network library[™] \(cuDNN\) 7.4.1](#)
- ▶ [NCCL 2.3.7](#) (optimized for [NVLink[™]](#))
- ▶ [OpenMPI 3.1.2](#)
- ▶ [TensorRT 5.0.2](#)

Driver Requirements

Release 18.12 is based on CUDA 10, which requires [NVIDIA Driver](#) release 410.xx. However, if you are running on Tesla (Tesla V100, Tesla P4, Tesla P40, or Tesla P100), you may use NVIDIA driver release 384. For more information, see [CUDA Compatibility and Upgrades](#).

GPU Requirements

Release 18.12 supports CUDA compute capability 6.0 and higher. This corresponds to GPUs in the Pascal, Volta, and Turing families. Specifically, for a list of GPUs that this compute capability corresponds to, see [CUDA GPUs](#). For additional support details, see [Deep Learning Frameworks Support Matrix](#).

Key Features and Enhancements

This NVCaffe release includes the following key features and enhancements.

- ▶ NVCaffe 0.17.2 container image version 18.12 is based on [Caffe Deep Learning Framework by the BVL](http://caffe.berkeleyvision.org/)C (http://caffe.berkeleyvision.org/).
- ▶ Ubuntu 16.04 with November 2018 updates

Known Issues

There are no known issues in this release.

Chapter 20.

NVCAFFE RELEASE 18.11

The NVIDIA container image of [Caffe](#), release 18.11, is available.

Contents of NVCAffe

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

The container also includes the following:

- ▶ [Ubuntu 16.04](#) including [Python 2.7](#)
- ▶ [NVIDIA CUDA 10.0.130](#) including [CUDA[®] Basic Linear Algebra Subroutines library[™] \(cuBLAS\) 10.0.130](#)
- ▶ [NVIDIA CUDA[®] Deep Neural Network library[™] \(cuDNN\) 7.4.1](#)
- ▶ [NCCL 2.3.7](#) (optimized for [NVLink[™]](#))
- ▶ [OpenMPI 3.1.2](#)
- ▶ [TensorRT 5.0.2](#)

Driver Requirements

Release 18.11 is based on CUDA 10, which requires [NVIDIA Driver](#) release 410.xx. However, if you are running on Tesla (Tesla V100, Tesla P4, Tesla P40, or Tesla P100), you may use NVIDIA driver release 384. For more information, see [CUDA Compatibility and Upgrades](#).

Key Features and Enhancements

This NVCAffe release includes the following key features and enhancements.

- ▶ NVCAffe 0.17.1 container image version 18.11 is based on [Caffe Deep Learning Framework](#) by the BVLC (<http://caffe.berkeleyvision.org/>).
- ▶ Latest version of [NCCL 2.3.7](#).
- ▶ Latest version of [NVIDIA cuDNN 7.4.1](#).

- ▶ Latest version of [TensorRT 5.0.2](#)
- ▶ Ubuntu 16.04 with October 2018 updates

Known Issues

There are no known issues in this release.

Chapter 21.

NVCAFFE RELEASE 18.10

The NVIDIA container image of [Caffe](#), release 18.10, is available.

Contents of NVCAffe

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

The container also includes the following:

- ▶ [Ubuntu 16.04](#) including [Python 2.7](#)
- ▶ [NVIDIA CUDA 10.0.130](#) including [CUDA[®] Basic Linear Algebra Subroutines library[™] \(cuBLAS\) 10.0.130](#)
- ▶ [NVIDIA CUDA[®] Deep Neural Network library[™] \(cuDNN\) 7.4.0](#)
- ▶ [NCCL 2.3.6](#) (optimized for [NVLink[™]](#))
- ▶ [OpenMPI 3.1.2](#)
- ▶ [TensorRT 5.0.0 RC](#)

Driver Requirements

Release 18.10 is based on CUDA 10, which requires [NVIDIA Driver](#) release 410.xx. However, if you are running on Tesla (Tesla V100, Tesla P4, Tesla P40, or Tesla P100), you may use NVIDIA driver release 384. For more information, see [CUDA Compatibility and Upgrades](#).

Key Features and Enhancements

This NVCAffe release includes the following key features and enhancements.

- ▶ NVCAffe 0.17.1 container image version 18.10 is based on [Caffe Deep Learning Framework](#) by the BVLC (<http://caffe.berkeleyvision.org/>).
- ▶ Latest version of [NCCL 2.3.6](#).
- ▶ Latest version of [OpenMPI 3.1.2](#).

- ▶ Ubuntu 16.04 with September 2018 updates

Known Issues

There are no known issues in this release.

Chapter 22.

NVCAFFE RELEASE 18.09

The NVIDIA container image of [Caffe](#), release 18.09, is available.

Contents of NVCaffe

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

The container also includes the following:

- ▶ Ubuntu 16.04 including Python 2.7
- ▶ NVIDIA CUDA 10.0.130 including CUDA[®] Basic Linear Algebra Subroutines library[™] (cuBLAS) 10.0.130
- ▶ NVIDIA CUDA[®] Deep Neural Network library[™] (cuDNN) 7.3.0
- ▶ NCCL 2.3.4 (optimized for NVLink[™])
- ▶ OpenMPI 2.0
- ▶ TensorRT 5.0.0 RC

Driver Requirements

Release 18.09 is based on CUDA 10, which requires [NVIDIA Driver](#) release 410.xx. However, if you are running on Tesla (Tesla V100, Tesla P4, Tesla P40, or Tesla P100), you may use NVIDIA driver release 384. For more information, see [CUDA Compatibility and Upgrades](#).

Key Features and Enhancements

This NVCaffe release includes the following key features and enhancements.

- ▶ NVCaffe 0.17.1 container image version 18.09 is based on [Caffe Deep Learning Framework](#) by the BVLC (<http://caffe.berkeleyvision.org/>).
- ▶ Latest version of cuDNN 7.3.0.

- ▶ Latest version of [CUDA 10.0.130](#) which includes support for DGX-2, Turing, and Jetson Xavier.
- ▶ Latest version of [cuBLAS 10.0.130](#).
- ▶ Latest version of [NCCL 2.3.4](#).
- ▶ Latest version of [TensorRT 5.0.0 RC](#).



All 18.09 containers inherit TensorRT 5.0.0 RC from the base container, however, some containers may not use TensorRT if there is no support for TensorRT in the given framework.

- ▶ Ubuntu 16.04 with August 2018 updates

Known Issues

There are no known issues in this release.

Chapter 23.

NVCAFFE RELEASE 18.08

The NVIDIA container image of [Caffe](#), release 18.08, is available.

Contents of NVCAffe

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

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[™]](#))
- ▶ [OpenMPI[™] 2.0](#)

Driver Requirements

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

Key Features and Enhancements

This NVCAffe release includes the following key features and enhancements.

- ▶ NVCAffe 0.17.1 container image version 18.08 is based on [Caffe Deep Learning Framework](#) by the BVLC (<http://caffe.berkeleyvision.org/>).
- ▶ Latest version of [cuDNN 7.2.1](#).
- ▶ Included support for [OpenMPI](#) version 2.0 in container.
- ▶ [Ubuntu 16.04](#) with July 2018 updates

Known Issues

There are no known issues in this release.

Chapter 24.

NVCAFFE RELEASE 18.07

The NVIDIA container image of [Caffe](#), release 18.07, is available.

Contents of NVCAffe

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

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 NVCAffe release includes the following key features and enhancements.

- ▶ NVCAffe 0.17.1 container image version 18.07 is based on [Caffe Deep Learning Framework by the BVLC](#) (<http://caffe.berkeleyvision.org/>).
- ▶ Latest version of [CUDA[®] Basic Linear Algebra Subroutines library[™] \(cuBLAS\) 9.0.425](#).
- ▶ Added multi-node support.
- ▶ Improved adaptive loss scaling stability.
- ▶ [Ubuntu 16.04](#) with June 2018 updates

Known Issues

There are no known issues in this release.

Chapter 25.

NVCAFFE RELEASE 18.06

The NVIDIA container image of [Caffe](#), release 18.06, is available.

Contents of NVCAffe

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

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 NVCAffe release includes the following key features and enhancements.

- ▶ NVCAffe 0.17.0 container image version 18.06 is based on [Caffe Deep Learning Framework by the BVLC](#) (<http://caffe.berkeleyvision.org/>).
- ▶ Improved adaptive scaling algorithm for mixed precision
- ▶ TensorRT integrated via new **TRTLayer** (see `examples/TensorRT/readme.txt` for details)
- ▶ Ubuntu 16.04 with May 2018 updates

Known Issues

There are no known issues in this release.

Chapter 26.

NVCAFFE RELEASE 18.05

The NVIDIA container image of [Caffe](#), release 18.05, is available.

Contents of NVCAffe

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

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 NVCAffe release includes the following key features and enhancements.

- ▶ NVCAffe 0.17.0 container image version 18.05 is based on [Caffe Deep Learning Framework by the BVLC](#) (<http://caffe.berkeleyvision.org/>).
- ▶ Fixed the [eltwise layer clobbers the first bottom blob](#) issue
- ▶ Improved SSD performance
- ▶ [Ubuntu 16.04](#) with April 2018 updates

Known Issues

There are no known issues in this release.

Chapter 27.

NVCAFFE RELEASE 18.04

The NVIDIA container image of [Caffe](#), release 18.04, is available.

Contents of NVCAffe

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

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 NVCAffe release includes the following key features and enhancements.

- ▶ NVCAffe 0.17.0 container image version 18.04 is based on [Caffe Deep Learning Framework by the BVLC](#) (<http://caffe.berkeleyvision.org/>).
- ▶ Incorporated all upstream GitHub Caffe code as of [PR 487](#).
- ▶ Added support for **AxpyLayer**
- ▶ Improved performance of **SplitLayer**
- ▶ Latest version of NCCL 2.1.15
- ▶ Ubuntu 16.04 with March 2018 updates

Known Issues

There are no known issues in this release.

Chapter 28.

NVCAFFE RELEASE 18.03

The NVIDIA container image of [Caffe](#), release 18.03, is available.

Contents of NVCAffe

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

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 NVCAffe release includes the following key features and enhancements.

- ▶ NVCAffe 0.16.6 container image version 18.03 is based on [Caffe Deep Learning Framework by the BVLC](#) (<http://caffe.berkeleyvision.org/>).
- ▶ `ImageDataLayer` parallelized to improve training performance when the dataset appears to have multiple files on a disk.
- ▶ Improved training accuracy. For more information, see: [NVIDIA Caffe: GitHub](#)
- ▶ Added support for RNN, Deconvolution, and SSD layers.
- ▶ Latest version of cuBLAS 9.0.333
- ▶ Latest version of cuDNN 7.1.1

- ▶ Ubuntu 16.04 with February 2018 updates

Known Issues

There are no known issues in this release.

Chapter 29.

NVCAFFE RELEASE 18.02

The NVIDIA container image of [Caffe](#), release 18.02, is available.

Contents of NVCAffe

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

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 NVCAffe release includes the following key features and enhancements.

- ▶ NVCAffe 0.16.5 container image version 18.02 is based on [Caffe Deep Learning Framework by the BVLC](#) (<http://caffe.berkeleyvision.org/>).
- ▶ Added the `ImageDataParameter` settings which belong to the `ImageDataLayer's` `LayerParameter` settings. More specifically:

- ▶ Added the threads parameter which is the number of Data Transformer threads per GPU executed by the ImageDataLayer.
- ▶ Added the cache parameter which ensures that the data is read once and put into the host memory.
- ▶ 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 30.

NVCAFFE RELEASE 18.01

The NVIDIA container image of [Caffe](#), release 18.01, is available.

Contents of NVCAffe

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

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 NVCAffe release includes the following key features and enhancements.

- ▶ NVCAffe 0.16.5 container image version 18.01 is based on [Caffe Deep Learning Framework by the BVLC \(<http://caffe.berkeleyvision.org/>\)](#).
- ▶ Added Layer-wise Adaptive Rate Control (LARC) parameters
- ▶ Added adaptive weight decay parameters
- ▶ Added adaptive momentum parameters
- ▶ Added adaptive global gradient scaling parameters
- ▶ Added advanced augmentation algorithms
- ▶ Implemented less memory consumption

- ▶ 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 31.

NVCAFFE RELEASE 17.12

The NVIDIA container image of [Caffe](#), release 17.12, is available.

NVCaffe container image version 17.12 is based on [NVCaffe 0.16.4](#).

Contents of NVCaffe

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

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 NVCaffe 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 32.

NVCAFFE RELEASE 17.11

The NVIDIA container image of [Caffe](#), release 17.11, is available.

NVCaffe container image version 17.11 is based on [NVCaffe 0.16.4](#).

Contents of NVCaffe

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

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 NVCaffe 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 33.

NVCAFFE RELEASE 17.10

The NVIDIA container image of [Caffe](#), release 17.10, is available.

NVCaffe container image version 17.10 is based on [NVCaffe 0.16.4](#).

Contents of NVCaffe

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

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 NVCaffe 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 34.

NVCAFFE RELEASE 17.09

The NVIDIA container image of [Caffe](#), release 17.09, is available.

NVCaffe container image version 17.09 is based on [NVCaffe 0.16.4](#).

Contents of NVCaffe

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

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 NVCaffe release includes the following key features and enhancements.

- ▶ Added the `img_rand_resize_lower` and `img_rand_resize_upper` parameters which specify that the variable-sized input image should be randomly resized.
- ▶ Added the `global_grad_scale` parameter which defines the constant C used to improve precision of back-propagation for float16 data storage.
- ▶ Added the `default_cudnn_math_override` parameter which sets the default `cudnn_math_override` value for every layer, if applicable.
- ▶ Added the `store_blobs_in_old_format` parameter which stores blobs in an old, less efficient BVLC-compatible format.

- ▶ Added the `cuda_math_override` parameter which sets the default `cudaMathType_t` value for all cuDNN-based computations in the current layer, if applicable, otherwise, it is ignored.
- ▶ Added the `ELUParameter` parameter stores parameters used by ELULayer.
- ▶ Added the `lambda` parameter which is used for Scaled Exponential Linear Unit (SELU™).
- ▶ 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 35.

NVCAFFE RELEASE 17.07

The NVIDIA container image of [Caffe](#), release 17.07, is available.

NVCaffe container image version 17.07 is based on [NVCaffe 0.16](#).

Contents of NVIDIA NVCaffe

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

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](#)
- ▶ [NCCL 2.0.3](#) (optimized for [NVLink[™]](#))

Key Features and Enhancements

This NVCaffe release includes the following key features and enhancements.

- ▶ Fixed a bug that was introduced in 17.05 where the log output from the `caffe time` command was excessively large.
- ▶ [Ubuntu 16.04](#) with June 2017 updates

Known Issues

There are no known issues in this release.

Chapter 36.

NVCAFFE RELEASE 17.06

The NVIDIA container image of [Caffe](#), release 17.06, is available.

NVCaffe container image version 17.06 is based on [NVCaffe 0.16](#).

Contents of NVCaffe

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

The container also includes the following:

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

Key Features and Enhancements

This NVCaffe release includes the following key features and enhancements.

- ▶ Added the shuffle parameter which is a data augmentation technique that improves accuracy of training your network.
- ▶ Added the cache parameter which ensures that the data is read once and put into the host memory.
- ▶ Added the min_lr parameter which ensures that the learning rate (lr) threshold is larger than 0.
- ▶ Improved determinism between multiple GPUs
- ▶ Ubuntu 16.04 with May 2017 updates

Known Issues

There are no known issues in this release.

Chapter 37.

NVCAFFE RELEASE 17.05

The NVIDIA container image of [Caffe](#), release 17.05, is available.

NVCaffe container image version 17.05 is based on [NVCaffe 0.16](#).

Contents of NVCaffe

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

The container also includes the following:

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

Key Features and Enhancements

This NVCaffe release includes the following key features and enhancements.

- ▶ Latest cuDNN release
- ▶ Added the `solver_data_type` parameter which is used for storing weights and history.
- ▶ Ubuntu 16.04 with April 2017 updates

Known Issues

DetectNet training can hang on NVCaffe 17.05.

Chapter 38.

NVCAFFE RELEASE 17.04

The NVIDIA container image of [Caffe](#), release 17.04, is available.

NVCaffe container image version 17.04 is based on [NVCaffe 0.16](#).

Contents of NVCaffe

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

The container also includes the following:

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

Key Features and Enhancements

This NVCaffe release includes the following key features and enhancements.

- ▶ Parallel testing on multi-GPU systems allows even memory distribution across all GPUs for better performance.
- ▶ All layers now support FP16 storage. This allows to avoid 32-16-32 bit conversions.
- ▶ Main feature: memory balancer between I/O buffers and convolution workspace. It maximizes memory utilization for better performance.
- ▶ Ubuntu 16.04 with March 2017 updates

Known Issues

There are no known issues in this release.

Chapter 39.

NVCAFFE RELEASE 17.03

The NVIDIA container image of [Caffe](#), release 17.03, is available.

NVCaffe container image version 17.03 is based on [NVCaffe 0.16](#).

Contents of NVCaffe

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

The container also includes the following:

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

Key Features and Enhancements

This NVCaffe release includes the following key features and enhancements.

- ▶ Supports overriding convolution algorithms for each layer and whole net
- ▶ Performance improvements in FP16 and Pseudo-FP16 modes
- ▶ Memory usage optimization allows for bigger batch sizes which improves performance. For example, you can fit ResNet50 to batch size of 64 per GPU on a GP100 board.
- ▶ All GPU kernels in weights update flow are now fused into one for each solver for better performance
- ▶ Ubuntu 16.04 with February 2017 updates

Known Issues

There are no known issues in this release.

Chapter 40.

NVCAFFE RELEASE 17.02

The NVIDIA container image of [Caffe](#), release 17.02, is available.

NVCaffe container image version 17.02 is based on [NVCaffe 0.16](#).

Contents of NVCaffe

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

The container also includes the following:

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

Key Features and Enhancements

This NVCaffe release includes the following key features and enhancements.

- ▶ Option to set the storage precision and match precision separately
- ▶ Supports ramp-up which allows you to choose a smaller learning rate when working with deep nets
- ▶ Data Reader is now multi-threaded for better performance, especially for I/O intensive nets
- ▶ Ubuntu 14.04 with January 2017 updates

Known Issues

There are no known issues in this release.

Chapter 41.

NVCAFFE RELEASE 17.01

The NVIDIA container image of [Caffe](#), release 17.01, is available.

NVCaffe container image version 17.01 is based on [NVCaffe 0.16](#).

Contents of NVCaffe

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

The container also includes the following:

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

Key Features and Enhancements

This NVCaffe release includes the following key features and enhancements.

- ▶ Parser is now multi-threaded for better performance, especially for I/O intensive nets
- ▶ Ubuntu 14.04 with December 2016 updates

Known Issues

There are no known issues in this release.

Chapter 42.

NVCAFFE RELEASE 16.12

The NVIDIA container image of [Caffe](#), release 16.12, is available.

NVCaffe container image version 16.12 is based on [NVCaffe 0.16](#).

Contents of NVCaffe

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

The container also includes the following:

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

Key Features and Enhancements

This NVCaffe release includes the following key features and enhancements.

- ▶ Batch Normalization Layer is backward compatible with BVLC
- ▶ Supports `pycaffe`, which makes the NVCaffe interfaces available for use through Python
- ▶ Ubuntu 14.04 with November 2016 updates

Known Issues

There are no known issues with 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, Triton 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.