

Developer Guide

Amazon Deep Learning AMIs



Amazon Deep Learning AMIs: Developer Guide

Table of Contents

What is the DLAMI?	1
About this guide	1
Prerequisites	1
Example use cases	1
Features	2
Preinstalled frameworks	2
Preinstalled GPU software	3
Model serving and visualization	3
DLAMI Release Notes	4
P6-B200 Supported DLAMI	4
P6-B200 Supported DLAMI	4
GPU Functionality Test	5
Base DLAMIs	8
X86	8
ARM64	62
Single-Framework DLAMIs	79
PyTorch DLAMIs	79
TensorFlow DLAMIs	123
Multi-Framework DLAMI	132
X86	132
Getting started	146
Choosing a DLAMI	146
CUDA Installations and Framework Bindings	147
Base	148
Conda	148
Architecture	150
OS	150
Choosing an instance	150
Pricing	152
Region Availability	152
GPU	152
CPU	153
Inferentia	153
Trainium	154

Setting up	156
Finding a DLAMI ID	156
Launching an instance	158
Connecting to an instance	160
Setting up Jupyter	160
Securing server	161
Starting server	162
Connecting client	162
Logging in	164
Cleaning up	165
Using a DLAMI	167
Conda DLAMI	167
Introduction to the Deep Learning AMI with Conda	167
Log in to Your DLAMI	168
Start the TensorFlow Environment	168
Switch to the PyTorch Python 3 Environment	169
Removing Environments	170
Base DLAMI	170
Using the Deep Learning Base AMI	170
Configuring CUDA Versions	171
Jupyter Notebooks	171
Navigating the Installed Tutorials	172
Switching Environments with Jupyter	172
Tutorials	173
Activating Frameworks	173
Elastic Fabric Adapter	176
GPU Monitoring and Optimization	189
Amazon Inferentia	199
ARM64 DLAMI	221
Inference	224
Model Serving	224
Upgrading Your DLAMI	229
DLAMI Upgrade	229
Software Updates	230
Release Notifications	231
Security	233

Data protection	234
Identity and access management	235
Authenticating with identities	235
Managing access using policies	238
IAM with Amazon EMR	240
Compliance validation	240
Resilience	241
Infrastructure security	242
Monitoring	242
Usage tracking	242
DLAMI Support Policy	243
DLAMI Support FAQs	243
What framework versions get security patches?	244
Which operating system get security patches?	244
What images does Amazon publish when new framework versions are released?	244
What images get new SageMaker AI/Amazon features?	244
How is current version defined in the Supported Frameworks table?	244
What if I am running a version that is not in the Supported table?	244
Do DLAMIs support previous patch versions of a Framework Version?	245
How can I find the latest patched image for a supported framework version?	245
How frequently are new images released?	245
Will my instance be patched in place while my workload is running?	245
What happens when a new patched or updated framework version is available?	246
Are dependencies updated without changing the framework version?	246
When does active support for my framework version end?	246
Will images with framework versions that are no longer actively maintained be	
patched?	247
How do I use an older framework version?	248
How do I stay up-to-date with support changes in frameworks and their versions?	248
Do I need a commercial license to use the Anaconda Repository?	248
DLAMI Support Policy Table	249
Supported Framework Versions	249
Supported Operating System Versions	249
Unsupported Framework Versions	249
Unsupported Operating System Versions	250
Unsupported DLAMI Release Notes Archive	251

Base	251
Single Framework	251
Multi Framework	253
Important changes	254
DLAMI NVIDIA driver change FAQs	254
What changed?	254
Why was this change required?	
Which DLAMIs did this change affect?	255
What does this mean for you?	256
Is there any loss of functionality with the newer DLAMIs?	. 256
Did this change affect Deep Learning Containers?	256
Related information	257
Deprecated features	258
-	

What is Amazon Deep Learning AMIs?

Amazon Deep Learning AMIs (DLAMI) provides customized machine images that you can use for deep learning in the cloud. The DLAMIs are available in most Amazon Web Services Regions for a variety of Amazon Elastic Compute Cloud (Amazon EC2) instance types, from a small CPUonly instance to the latest high-powered multi-GPU instances. The DLAMIs come preconfigured with <u>NVIDIA CUDA</u> and <u>NVIDIA cuDNN</u> and the latest releases of the most popular deep learning frameworks.

About this guide

The content in the can help you launch and use the DLAMIs. The guide covers several common deep learning use cases, for both training and inference. It also covers how to choose the right AMI for your purpose and the kind of instances that you might prefer.

Additionally, the DLAMIs include several tutorials that their supported frameworks provide. This guide can show you how to activate each framework and find the appropriate tutorials to get started. It also has tutorials on distributed training, debugging, using Amazon Inferentia and Amazon Trainium, and other key concepts. For instructions on how to set up a Jupyter notebook server to run the tutorials in your browser, see <u>Setting up a Jupyter Notebook server on a DLAMI instance</u>.

Prerequisites

To successfully run the DLAMIs, we recommend that you be familiar with command line tools and basic Python.

Example DLAMI use cases

The following are examples of some common use cases for Amazon Deep Learning AMIs (DLAMI).

Learning about deep learning – DLAMI is a great choice for learning or teaching machine learning and deep learning frameworks. The DLAMIs take away the headache from troubleshooting the installations of each framework and getting them to play along on the same computer. The DLAMIs include a Jupyter notebook and make it easy to run the tutorials that the frameworks provide for people new to machine learning and deep learning. **App development** – If you're an app developer who's interested in using deep learning to make your apps utilize the latest advances in AI, then DLAMI is the perfect test bed for you. Each framework comes with tutorials on how to get started with deep learning, and many of them have model zoos that make it easy to try out deep learning without having to create the neural networks yourself or to do any of the model training. Some examples show you how to build an image detection application in just a few minutes, or how to build a speech recognition app for your own chatbot.

Machine learning and data analytics – If you're a data scientist or you're interested in processing your data with deep learning, then you'll find that many of the frameworks have support for R and Spark. You will find tutorials on how to do simple regressions, all the way up to building scalable data processing systems for personalization and predictions systems.

Research – If you're a researcher who wants to try out a new framework, test a new model, or train new models, then DLAMI and Amazon capabilities for scale can alleviate the pain of tedious installations and management of multiple training nodes.

i Note

While your initial choice might be to upgrade your instance type to a larger instance with more GPUs (up to 8), you can also scale horizontally by creating a cluster of DLAMI instances. Check out <u>Related information about DLAMI</u> for more information on cluster builds.

Features of DLAMI

The features of Amazon Deep Learning AMIs (DLAMI) include preinstalled deep learning frameworks, GPU software, model servers, and model visualization tools.

Preinstalled frameworks

There are currently two primary flavors of DLAMI with other variations related to the operating system (OS) and software versions:

- <u>Deep Learning AMI with Conda</u> Frameworks installed separately using conda packages and separate Python environments.
- Deep Learning Base AMI No frameworks installed; only NVIDIA CUDA and other dependencies.

The Deep Learning AMI with Conda uses conda environments to isolate each framework, so you can switch between them at will and not worry about their dependencies conflicting. The Deep Learning AMI with Conda supports the following frameworks:

- PyTorch
- TensorFlow 2

i Note

DLAMI no longer supports the following deep learning frameworks: Apache MXNet, Microsoft Cognitive Toolkit (CNTK), Caffe, Caffe2, Theano, Chainer, and Keras.

Preinstalled GPU software

Even if you use a CPU-only instance, the DLAMIs will have <u>NVIDIA CUDA</u> and <u>NVIDIA cuDNN</u>. The installed software is the same regardless of the instance type. Keep in mind that GPU-specific tools work only on an instance that has at least one GPU. For more information about instance types, see <u>Choosing a DLAMI instance type</u>.

For more information about CUDA, see CUDA Installations and Framework Bindings.

Model serving and visualization

Deep Learning AMI with Conda comes preinstalled with model servers for TensorFlow, as well as TensorBoard for model visualizations. For more information, see <u>TensorFlow Serving</u>.

Deep Learning AMIs Release Notes

Here you can find detailed release notes for all currently supported Amazon Deep Learning AMIs (DLAMI) options.

For release notes for DLAMI frameworks that we no longer support, see the **Unsupported Framework Release Notes Archive** section of the <u>DLAMI Framework Support Policy</u> page.

1 Note

The Amazon Deep Learning AMIs have a nightly release cadence for security patches. We don't include these incremental security patches in the release notes.

Release Notes

- P6-B200 Supported DLAMI
- <u>Release Notes for Base DLAMIs</u>
- Release Notes for Single Framework DLAMIs
- <u>Release Notes for Multi-Framework DLAMIs</u>

P6-B200 Supported DLAMI

Below are the detailed requirements for running DLAMI on

P6-B200 Supported DLAMIs

The Following DLAMI support P6-B200 instances:

- Amazon Deep Learning Base AMI (Amazon Linux 2023)
- Amazon Deep Learning Base AMI (Ubuntu 24.04)
- Amazon Deep Learning Base AMI (Ubuntu 22.04)

These DLAMI contain the following software required to operate P6-B200 instances:

Software	Minimum Version Requirement
Nvidia CUDA Toolkit	12.8
Nvidia Driver	R570
NVLINK 5	R570
Linux Kernel	6.1
Elastic Fabric Adapter (EFA)	1.41.0
Amazon OFI NCCL Plugin	1.15.0

Confirm GPU Functionality

To confirm functional GPUs:

1. Run the following Nvidia GPU Device Query Test

```
$ /usr/local/cuda/extras/demo_suite/deviceQuery
```

2. Confirm the Following output from the Device Query Run:

```
$ /usr/local/cuda/extras/demo_suite/deviceQuery
/usr/local/cuda/extras/demo_suite/deviceQuery Starting...
CUDA Device Query (Runtime API)
Detected 8 CUDA Capable device(s)
...
deviceQuery, CUDA Driver = CUDART, CUDA Driver Version = 12.8, CUDA Runtime Version
= 12.8, NumDevs = 8, Device0 = NVIDIA B200, Device1 = NVIDIA B200, Device2 =
NVIDIA B200, Device3 = NVIDIA B200, Device4 = NVIDIA B200, Device5 = NVIDIA B200,
Device6 = NVIDIA B200, Device7 = NVIDIA B200
Result = PASS
```

To confirm functional NVIDIA Driver:

1. Run the Nvidia System Management Interface

```
$ nvidia-smi
```

2. Confirm the Following output from the System Management Interface

```
_____
+
NVIDIA-SMI 570.133.20 Driver Version: 570.133.20 CUDA Version:
12.8
    +----+
GPU Name
             Persistence-M | Bus-Id Disp.A | Volatile
Uncorr. ECC |
| Fan Temp Perf Pwr:Usage/Cap | Memory-Usage | GPU-Util
Compute M.
                     L
MIG M. |
Off | 0000000:51:00.0 Off |
0 NVIDIA B200
  0
          145W / 1000W | 0MiB / 183359MiB |
| N/A 32C
       P0
                                      0%
Default |
                     L
Disabled |
-----
           ----+
                   Off |
                       00000000:52:00.0 Off |
1 NVIDIA B200
 0
       P0 140W / 1000W | 0MiB / 183359MiB |
| N/A 30C
                                      0%
Default |
                     Disabled |
+----+
                   Off | 0000000:62:00.0 Off |
2 NVIDIA B200
  0
       P0 139W / 1000W | 0MiB / 183359MiB |
| N/A 31C
                                      0%
Default |
                     Disabled |
```

GPU Functionality Test

3 NVIDIA H 0	B200	Off	00000000:63:00.0 Off	I
N/A 29C	P0	139W / 1000W	0MiB / 183359MiB	0%
isabled		I		I
		+-		
4 NVIDIA 8 0		Off	00000000:75:00.0 Off	I
N/A 31C efault	PØ	141W / 1000W	0MiB / 183359MiB	0%
isabled		I		I
		+-		
5 NVIDIA I 0		Off	00000000:76:00.0 Off	I
N/A 31C Default	P0	141W / 1000W	0MiB / 183359MiB	0%
Disabled				
		+-		
6 NVIDIA H 0	B200	0ff	00000000:86:00.0 Off	I
N/A 32C Default	P0	141W / 1000W	0MiB / 183359MiB	0%
)isabled		I		I
	+	+-		
7 NVIDIA 0	B200	Off	00000000:87:00.0 Off	I
N/A 30C Default	PØ	138W / 1000W	0MiB / 183359MiB	0%
Disabled		I		I
	+	+-		

+

_ _

++				
Processes:				
 GPU GI CI	PID	Туре	Process name	GPU
Memory		51		
ID ID				
Usage I				
ı ====================================		=======		
No running processes fo	ound			
+				
+				

If you experience any issues with P6-B200 instances, please reach out to Amazon Support.

Release Notes for Base DLAMIs

X86 Base DLAMI Release Notes

- X86 Base DLAMI Release Notes
- ARM64 Base DLAMI Release Notes

X86 Base DLAMI Release Notes

Below are the release notes for X86 Base DLAMI:

GPU

- Amazon Deep Learning Base AMI (Amazon Linux 2023) (Supports P6-B200)
- Amazon Deep Learning Base AMI (Ubuntu 24.04) (Supports P6-B200)
- Amazon Deep Learning Base AMI (Ubuntu 22.04) (Supports P6-B200)
- Amazon Deep Learning Base AMI (Amazon Linux 2)

Qualcomm

• Amazon Deep Learning Base Qualcomm AMI (Amazon Linux 2)

Amazon Neuron

• Refer to the Neuron DLAMI User Guide.

Amazon Deep Learning Base GPU AMI (Amazon Linux 2023)

For help getting started, see Getting started with DLAMI.

AMI name format

• Deep Learning Base OSS Nvidia Driver GPU AMI (Amazon Linux 2023) \${YYYY-MM-DD}

Supported EC2 Instances

- Please refer to Important changes to DLAMI
- Deep Learning with OSS Nvidia Driver supports G4dn, G5, G6, Gr6, G6e, P4d, P4de, P5, P5e, P5en, P6-B200

The AMI includes the following:

- Supported Amazon Service: Amazon EC2
- Operating System: Amazon Linux 2023
- Compute Architecture: x86
- Latest available version is installed for the following packages:
 - Linux Kernel: 6.1
 - FSx Lustre
 - NVIDIA GDS
 - Docker
 - Amazon CLI v2 at /usr/local/bin/aws2 and Amazon CLI v1 at /usr/bin/aws
 - NVIDIA DCGM
 - Nvidia container toolkit:
 - Version command: nvidia-container-cli -V
 - Nvidia-docker2:
 - Version command: nvidia-docker version
- NVIDIA Driver: 570.133.20

• NVIDIA CUDA12.4-12.6 and 12.8 stack:

- CUDA, NCCL and cuDDN installation directories: /usr/local/cuda-xx.x/
 - Example: /usr/local/cuda-12.8/ , /usr/local/cuda-12.8/
- Compiled NCCL Version: 2.26.5
- Default CUDA: 12.8
 - PATH /usr/local/cuda points to CUDA 12.8
 - Updated below env vars:
 - LD_LIBRARY_PATH to have /usr/local/cuda-12.8/lib:/usr/local/cuda-12.8/lib64:/usr/local/cuda-12.8:/usr/local/cuda-12.4/targets/x86_64-linux/lib
 - PATH to have /usr/local/cuda-12.8/bin/:/usr/local/cuda-12.8/include/
 - For any different CUDA version, please update LD_LIBRARY_PATH accordingly.
- EFA installer: 1.40.0
- Nvidia GDRCopy: 2.5.1
- Amazon OFI NCCL plugin comes with EFA installer
 - Paths /opt/amazon/ofi-nccl/lib and /opt/amazon/ofi-nccl/efa are added to LD_LIBRARY_PATH.
- Amazon CLI v2 at /usr/local/bin/aws2 and Amazon CLI v1 at /usr/bin/aws
- EBS volume type: gp3
- Python: /usr/bin/python3.9
- NVMe Instance Store Location (on Supported EC2 instances): /opt/dlami/nvme
- Query AMI-ID with SSM Parameter (example Region is us-east-1):
 - OSS Nvidia Driver:

```
SSM_PARAMETER=base-oss-nvidia-driver-gpu-amazon-linux-2023/latest/ami-id \
    aws ssm get-parameter --region us-east-1 \
    --name /aws/service/deeplearning/ami/x86_64/$SSM_PARAMETER \
    --query "Parameter.Value" \
    --output text
```

- Query AMI-ID with AWSCLI (example Region is us-east-1):
 - OSS Nvidia Driver:

```
aws ec2 describe-images --region us-east-1 \
    --owners amazon \
```

```
--filters 'Name=name,Values=Deep Learning Base OSS Nvidia Driver GPU AMI
(Amazon Linux 2023) ??????' 'Name=state,Values=available' \
    --query 'reverse(sort_by(Images, &CreationDate))[:1].ImageId' \
    --output text
```

Notices

NVIDIA Container Toolkit 1.17.4

In Container Toolkit version 1.17.4 the mounting of CUDA compat libraries is now disabled. In order to ensure compatibility with multiple CUDA versions on container workflows, please ensure you update your LD_LIBRARY_PATH to include your CUDA compatibility libraries as shown in the <u>If you</u> <u>use a CUDA compatibility layer</u> tutorial.

Support policy

These AMIs Components of this AMI like CUDA versions may be removed and changed based on <u>framework support policy</u> or to optimize performance for <u>deep learning containers</u> or to reduce AMI size in a future release, without prior notice. We remove CUDA versions from AMIs if they are not used by any supported framework version.

P6-B200 instances

P6-B200 instances contain 8 network interface cards, and can be launched using the following Amazon CLI command:

```
"NetworkCardIndex=5,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa"
\
     "NetworkCardIndex=6,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa"
     "NetworkCardIndex=7,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa"
```

P5en instances

P5en instances contain 16 network interface cards, and can be launched using the following Amazon CLI command:

```
aws ec2 run-instances --region $REGION \
    --instance-type $INSTANCETYPE \
    --image-id $AMI --key-name $KEYNAME \
    --iam-instance-profile "Name=dlami-builder" \
    --tag-specifications "ResourceType=instance,Tags=[{Key=Name,Value=$TAG}]" \
    --network-interfaces "NetworkCardIndex=0, DeviceIndex=0, Groups=$SG, SubnetId=
$SUBNET,InterfaceType=efa" \
      "NetworkCardIndex=1,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa"
/
      "NetworkCardIndex=2,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa"
/
      "NetworkCardIndex=3,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa"
\mathbf{X}
      "NetworkCardIndex=4,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa"
/
      . . .
      "NetworkCardIndex=15, DeviceIndex=1, Groups=$SG, SubnetId=$SUBNET, InterfaceType=efa"
```

P5/P5e instances

P5 and P5e instances contain 32 network interface cards, and can be launched using the following Amazon CLI command:

```
aws ec2 run-instances --region $REGION \
    --instance-type $INSTANCETYPE \
    --image-id $AMI --key-name $KEYNAME \
    --iam-instance-profile "Name=dlami-builder" \
    --tag-specifications "ResourceType=instance,Tags=[{Key=Name,Value=$TAG}]" \
    --network-interfaces "NetworkCardIndex=0,DeviceIndex=0,Groups=$SG,SubnetId=
$SUBNET,InterfaceType=efa" \
    "NetworkCardIndex=1,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa" \
```

```
"NetworkCardIndex=2,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa"
    "NetworkCardIndex=3,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa"
    "NetworkCardIndex=4,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa"
    ...
    "NetworkCardIndex=31,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa"
```

Kernel

• Kernel version is pinned using command:

sudo dnf versionlock kernel*

• We recommend that users avoid updating their kernel version (unless due to a security patch) to ensure compatibility with installed drivers and package versions. If users still wish to update, they can run the following commands to unpin their kernel versions:

```
sudo dnf versionlock delete kernel*
sudo dnf update -y
```

• For each new version of DLAMI, latest available compatible kernel is used.

Release Date: 2025-05-15

AMI name: Deep Learning Base OSS Nvidia Driver GPU AMI (Amazon Linux 2023) 20250515

Added

Added support for <u>P6-B200 EC2 instances</u>

Updated

- Upgraded EFA Installer from version 1.38.1 to 1.40.0
- Upgraded GDRCopy from version 2.4 to 2.5
- Upgraded Amazon OFI NCCL Plugin from version 1.13.0-aws to 1.14.2-aws
- Updated compiled NCCL Version from version 2.25.1 to 2.26.5
- Updated default CUDA version from version 12.6 to 12.8

• Updated Nvidia DCGM version from 3.3.9 to 4.4.3

Release Date: 2025-04-22

AMI name: Deep Learning Base OSS Nvidia Driver GPU AMI (Amazon Linux 2023) 20250421

Updated

 Upgraded Nvidia driver from version 570.124.06 to 570.133.20 to address CVEs present in the NVIDIA GPU Display Driver Security Bulletin for April 2025

Release Date: 2025-03-31

AMI name: Deep Learning Base OSS Nvidia Driver GPU AMI (Amazon Linux 2023) 20250328

Added

Added support for NVIDIA GPU Direct Storage (GDS)

Release Date: 2025-02-17

AMI name: Deep Learning Base OSS Nvidia Driver GPU AMI (Amazon Linux 2023) 20250215

Updated

- Updated NVIDIA Container Toolkit from version 1.17.3 to version 1.17.4
 - Please see the release notes page here for more information: <u>https://github.com/NVIDIA/</u> <u>nvidia-container-toolkit/releases/tag/v1.17.4</u>
 - In Container Toolkit version 1.17.4, the mounting of CUDA compat libraries is now disabled. In order to ensure compatibility with multiple CUDA versions on container workflows, please ensure you update your LD_LIBRARY_PATH to include your CUDA compatibility libraries as shown in the If you use a CUDA compatibility layer tutorial.

Removed

 Removed user space libraries cuobj and nvdisasm provided by <u>NVIDIA CUDA toolkit</u> to address CVEs present in the NVIDIA CUDA Toolkit Security Bulletin for February 18, 2025

Release Date: 2025-02-05

AMI name: Deep Learning Base OSS Nvidia Driver GPU AMI (Amazon Linux 2023) 20250205

Added

- Added CUDA toolkit version 12.6 in directory /usr/local/cuda-12.6
- Added support for G5 EC2 Instances

Removed

- CUDA versions 12.1 and 12.2 has been removed from this DLAMI. Customers who require these CUDA toolkit versions can install them directly from NVIDIA using the link below
 - https://developer.nvidia.com/cuda-toolkit-archive

Release Date: 2025-02-03

AMI name: Deep Learning Base OSS Nvidia Driver GPU AMI (Amazon Linux 2023) 20250131

Updated

- Upgraded EFA version from 1.37.0 to 1.38.0
 - EFA now bundles the Amazon OFI NCCL plugin, which can now be found in /opt/amazon/ofinccl rather than the original /opt/aws-ofi-nccl/. If updating your LD_LIBRARY_PATH variable, please ensure that you modify your OFI NCCL location properly.
- Upgraded Nvidia Container Toolkit from 1.17.3 to 1.17.4

Release Date: 2025-01-08

AMI name: Deep Learning Base OSS Nvidia Driver GPU AMI (Amazon Linux 2023) 20250107

Updated

Added support for <u>G4dn instances</u>

Release Date: 2024-12-09

AMI name: Deep Learning Base OSS Nvidia Driver GPU AMI (Amazon Linux 2023) 20241206

Updated

• Upgraded Nvidia Container Toolkit from version 1.17.0 to 1.17.3

Release Date: 2024-11-21

AMI name: Deep Learning Base OSS Nvidia Driver GPU AMI (Amazon Linux 2023) 20241121

Added

• Added support for P5en EC2 Instances.

Updated

- Upgraded EFA Installer from version 1.35.0 to 1.37.0
- Upgrade Amazon OFI NCCL Plugin from version 1.121-aws to 1.13.0-aws

Release Date: 2024-10-30

AMI name: Deep Learning Base OSS Nvidia Driver GPU AMI (Amazon Linux 2023) 20241030

Added

• Initial release of the Deep Learning Base OSS DLAMI for Amazon Linux 2023

Known Issues

This DLAMI does not support G4dn and G5 EC2 instances at this time. Amazon is aware of an incompatibility that may result in CUDA initialization failures, affecting both G4dn and G5 instance families when using the open source NVIDIA drivers together with a Linux kernel version 6.1 or newer. This issue affects Linux distributions such as Amazon Linux 2023, Ubuntu 22.04 or newer, or SUSE Linux Enterprise Server 15 SP6 or newer, among others.

Amazon Deep Learning Base GPU AMI (Ubuntu 24.04)

For help getting started, see <u>Getting started with DLAMI</u>.

AMI name format

• Deep Learning Base OSS Nvidia Driver GPU AMI (Ubuntu 24.04) \${YYYY-MM-DD}

Supported EC2 instances

- Please refer to Important changes to DLAMI.
- Deep Learning with OSS Nvidia Driver supports G4dn, G5, G6, Gr6, G6e, P4d, P4de, P5, P5e, P5en, P6-B200.

The AMI includes the following:

- Supported Amazon Service: Amazon EC2
- Operating System: Ubuntu 24.04
- **Compute Architecture**: x86
- Latest available version is installed for the following packages:
 - Linux Kernel: 6.8
 - FSx Lustre
 - Docker
 - Amazon CLI v2 at /usr/bin/aws
 - NVIDIA DCGM
 - Nvidia container toolkit:
 - Version command: nvidia-container-cli -V
 - Nvidia-docker2:
 - Version command: nvidia-docker version
- NVIDIA Driver: 570.133.20
- NVIDIA CUDA12.6 and 12.8 stack:
 - CUDA, NCCL and cuDDN installation directories: /usr/local/cuda-xx.x/
 - Example: /usr/local/cuda-12.8/ , /usr/local/cuda-12.8/
 - Compiled NCCL Version: 2.25.1
 - Default CUDA: 12.8
 - PATH /usr/local/cuda points to CUDA 12.8
 - Updated below env vars:

- LD_LIBRARY_PATH to have /usr/local/cuda-12.8/lib:/usr/local/cuda-12.8/lib64:/usr/local/cuda-12.8/targets/sbsa-linux/lib:/usr/local/cuda-12.8/nvvm/lib64:/usr/local/cuda-12.8/extras/CUPTI/lib64
- PATH to have /usr/local/cuda-12.8/bin/:/usr/local/cuda-12.8/include/
- For any different CUDA version, please update LD_LIBRARY_PATH accordingly.
- EFA installer: 1.40.0
- Nvidia GDRCopy: 2.5.1
- Amazon OFI NCCL plugin comes with EFA installer
 - Paths /opt/amazon/ofi-nccl/lib and /opt/amazon/ofi-nccl/efa are added to LD_LIBRARY_PATH.
- Amazon CLI v2 at /usr/bin/aws
- EBS volume type: gp3
- Python: /usr/bin/python3.12
- NVMe Instance Store Location (on Supported EC2 instances): /opt/dlami/nvme
- Query AMI-ID with SSM Parameter (example Region is us-east-1):
 - OSS Nvidia Driver:

```
SSM_PARAMETER=base-oss-nvidia-driver-gpu-ubuntu-24.04/latest/ami-id \
    aws ssm get-parameter --region us-east-1 \
    --name /aws/service/deeplearning/ami/x86_64/$SSM_PARAMETER \
    --query "Parameter.Value" \
    --output text
```

- Query AMI-ID with AWSCLI (example Region is us-east-1):
 - OSS Nvidia Driver:

```
aws ec2 describe-images --region us-east-1 \
    --owners amazon \
    --filters 'Name=name,Values=Deep Learning Base OSS Nvidia Driver GPU AMI
 (Ubuntu 24.04) ??????' 'Name=state,Values=available' \
    --query 'reverse(sort_by(Images, &CreationDate))[:1].ImageId' \
    --output text
```

Notices

Support policy

These AMIs Components of this AMI like CUDA versions may be removed and changed based on <u>framework support policy</u> or to optimize performance for <u>deep learning containers</u> or to reduce AMI size in a future release, without prior notice. We remove CUDA versions from AMIs if they are not used by any supported framework version.

EC2 instance with multiple network cards

- Many instances types that support EFA also have multiple network cards.
- DeviceIndex is unique to each network card, and must be a non-negative integer less than the limit of ENIs per NetworkCard. On P5, the number of ENIs per NetworkCard is 2, meaning that the only valid values for DeviceIndex is 0 or 1.
 - For the primary network interface (network card index 0, device index 0), create an EFA (EFA with ENA) interface. You can't use an EFA-only network interface as the primary network interface.
 - For each additional network interface, use the next unused network card index, device index 1, and either an EFA (EFA with ENA) or EFA-only network interface, depending on your use case, such as ENA bandwidth requirements or IP address space. For example use cases, see EFA configuration for a P5 instances.
 - For more information, see the EFA Guide here.

P6-B200 instances

P6-B200 instances contain 8 network interface cards, and can be launched using the following Amazon CLI command:

```
aws ec2 run-instances --region $REGION \
    --instance-type $INSTANCETYPE \
    --image-id $AMI --key-name $KEYNAME \
    --iam-instance-profile "Name=dlami-builder" \
    --tag-specifications "ResourceType=instance,Tags=[{Key=Name,Value=$TAG}]" \
    --network-interfaces "NetworkCardIndex=0,DeviceIndex=0,Groups=$SG,SubnetId=
$SUBNET,InterfaceType=efa" \
    "NetworkCardIndex=1,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa" \
```



P5en instances

P5en contain 16 network interface cards, and can be launched using the following Amazon CLI command:

```
aws ec2 run-instances --region $REGION \
    --instance-type $INSTANCETYPE \
    --image-id $AMI --key-name $KEYNAME \
    --iam-instance-profile "Name=dlami-builder" \
    --tag-specifications "ResourceType=instance,Tags=[{Key=Name,Value=$TAG}]" \
    --network-interfaces "NetworkCardIndex=0,DeviceIndex=0,Groups=$SG,SubnetId=
$SUBNET,InterfaceType=efa" \
      "NetworkCardIndex=1,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa"
/
      "NetworkCardIndex=2,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa"
/
      "NetworkCardIndex=3,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa"
/
      "NetworkCardIndex=4,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa"
/
      . . .
      "NetworkCardIndex=15,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa"
```

P5/P5e instances

P5 and P5e instances contain 32 network interface cards, and can be launched using the following Amazon CLI command:

```
aws ec2 run-instances --region $REGION \
--instance-type $INSTANCETYPE \
```

```
--image-id $AMI --key-name $KEYNAME \
--iam-instance-profile "Name=dlami-builder" \
--tag-specifications "ResourceType=instance,Tags=[{Key=Name,Value=$TAG}]" \
--network-interfaces "NetworkCardIndex=0,DeviceIndex=0,Groups=$SG,SubnetId=
$SUBNET,InterfaceType=efa" \
    "NetworkCardIndex=1,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa" \
    "NetworkCardIndex=3,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa" \
    "NetworkCardIndex=4,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa" \
    "NetworkCardIndex=4,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa" \
    "NetworkCardIndex=3,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa" \
    "NetworkCardIndex=3,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa" \
    "NetworkCardIndex=31,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa" \
    "NetworkCardIndex=31,DeviceIndex=1,Groups=$SG,SubnetIndex=30,SubnetId=$SUBNET,InterfaceType=efa" \
    "Ne
```

Kernel

• Kernel version is pinned using command:

```
echo linux-aws hold | sudo dpkg —set-selections
echo linux-headers-aws hold | sudo dpkg —set-selections
echo linux-image-aws hold | sudo dpkg —set-selections
```

 We recommend that users avoid updating their kernel version (unless due to a security patch) to ensure compatibility with installed drivers and package versions. If users still wish to update, they can run the following commands to unpin their kernel versions:

```
echo linux-aws install | sudo dpkg -set-selections
echo linux-headers-aws install | sudo dpkg -set-selections
echo linux-image-aws install | sudo dpkg -set-selections
```

• For each new version of DLAMI, latest available compatible kernel is used.

Release Date: 2025-05-22

AMI name: Deep Learning Base OSS Nvidia Driver GPU AMI (Ubuntu 24.04) 20250522

Added

Added support for <u>P6-B200 EC2 instances</u>

Updated

Upgraded EFA Installer from version 1.40.0 to 1.41.0

- Updated compiled NCCL Version from version 2.25.1 to 2.26.5
- Updated Nvidia DCGM version from 3.3.9 to 4.4.3

Release Date: 2025-05-13

AMI name: Deep Learning Base OSS Nvidia Driver GPU AMI (Ubuntu 24.04) 20250513

Added

• Initial release of the Deep Learning Base OSS DLAMI for Ubuntu 24.04

Amazon Deep Learning Base GPU AMI (Ubuntu 22.04)

For help getting started, see Getting started with DLAMI.

AMI name format

• Deep Learning Base OSS Nvidia Driver GPU AMI (Ubuntu 22.04) \${YYYY-MM-DD}

Supported EC2 instances

- Please refer to Important changes to DLAMI.
- Deep Learning with OSS Nvidia Driver supports G4dn, G5, G6, Gr6, G6e, P4d, P4de, P5, P5e, P6-B200.

The AMI includes the following:

- Supported Amazon Service: Amazon EC2
- Operating System: Ubuntu 22.04
- Compute Architecture: x86
- Latest available version is installed for the following packages:
 - Linux Kernel: 6.8
 - FSx Lustre
 - Docker
 - Amazon CLI v2 at /usr/local/bin/aws2 and Amazon CLI v1 at /usr/bin/aws
 - NVIDIA DCGM

- Nvidia container toolkit:
 - Version command: nvidia-container-cli -V
- Nvidia-docker2:
 - Version command: nvidia-docker version
- NVIDIA Driver: 570.133.20
- NVIDIA CUDA12.4-12.6 and 12.8 stack:
 - CUDA, NCCL and cuDDN installation directories: /usr/local/cuda-xx.x/
 - Example: /usr/local/cuda-12.8/ , /usr/local/cuda-12.8/
 - Compiled NCCL Version: 2.26.5
 - Default CUDA: 12.8
 - PATH /usr/local/cuda points to CUDA 12.8
 - Updated below env vars:
 - LD_LIBRARY_PATH to have /usr/local/cuda-12.8/lib:/usr/local/cuda-12.8/lib64:/usr/local/cuda-12.8:/usr/local/cuda-12.8/targets/x86_64-linux/lib:/usr/local/cuda-12.8/extras/CUPTI/lib64
 - PATH to have /usr/local/cuda-12.8/bin/:/usr/local/cuda-12.8/include/
 - For any different CUDA version, please update LD_LIBRARY_PATH accordingly.
- EFA installer: 1.40.0
- Nvidia GDRCopy: 2.5.1
- Amazon OFI NCCL plugin comes with EFA installer
 - Paths /opt/amazon/ofi-nccl/lib and /opt/amazon/ofi-nccl/efa are added to LD_LIBRARY_PATH.
- Amazon CLI v2 at /usr/local/bin/aws2 and Amazon CLI v1 at /usr/bin/aws
- EBS volume type: gp3
- **Python**: /usr/bin/python3.10
- NVMe Instance Store Location (on <u>Supported EC2 instances</u>): /opt/dlami/nvme
- Query AMI-ID with SSM Parameter (example Region is us-east-1):
 - OSS Nvidia Driver:

SSM_PARAMETER=base-oss-nvidia-driver-gpu-ubuntu-22.04/latest/ami-id \

```
--name /aws/service/deeplearning/ami/x86_64/$SSM_PARAMETER \
--query "Parameter.Value" \
--output text
```

- Query AMI-ID with AWSCLI (example Region is us-east-1):
 - OSS Nvidia Driver:

```
aws ec2 describe-images --region us-east-1 \
    --owners amazon \
    --filters 'Name=name,Values=Deep Learning Base OSS Nvidia Driver GPU AMI
 (Ubuntu 22.04) ???????' 'Name=state,Values=available' \
    --query 'reverse(sort_by(Images, &CreationDate))[:1].ImageId' \
    --output text
```

Notices

NVIDIA Container Toolkit 1.17.4

In Container Toolkit version 1.17.4 the mounting of CUDA compat libraries is now disabled. In order to ensure compatibility with multiple CUDA versions on container workflows, please ensure you update your LD_LIBRARY_PATH to include your CUDA compatibility libraries as shown in the <u>If you</u> <u>use a CUDA compatibility layer</u> tutorial.

EFA updates from 1.37 to 1.38 (Release on 2025-01-31)

EFA now bundles the Amazon OFI NCCL plugin, which can now be found in /opt/amazon/ofi-nccl rather than the original /opt/aws-ofi-nccl/. If updating your LD_LIBRARY_PATH variable, please ensure that you modify your OFI NCCL location properly.

Multi ENI support

- Ubuntu 22.04 automatically sets up and configures source routing on multiple NICss using cloudinit on its initial boot. If your workflow includes attaching/detaching your ENIs while an instance is stopped, an additional configuration must be added to the cloud-init user data to ensure proper configuration of the NIC's during these events. A sample of the cloud config is provided below.
- Please reference this Canonical documentation here for more information on how to configure the cloud config for your instances - <u>https://documentation.ubuntu.com/aws/en/latest/aws-</u> how-to/instances/automatically-setup-multiple-nics/

```
#cloud-config
# apply network config on every boot and hotplug event
updates:
    network:
    when: ['boot', 'hotplug']
```

Support policy

These AMIs Components of this AMI like CUDA versions may be removed and changed based on <u>framework support policy</u> or to optimize performance for <u>deep learning containers</u> or to reduce AMI size in a future release , without prior notice. We remove CUDA versions from AMIs if they are not used by any supported framework version.

EC2 instances with multiple network cards

- Many instances types that support EFA also have multiple network cards.
- DeviceIndex is unique to each network card, and must be a non-negative integer less than the limit of ENIs per NetworkCard. On P5, the number of ENIs per NetworkCard is 2, meaning that the only valid values for DeviceIndex is 0 or 1.
 - For the primary network interface (network card index 0, device index 0), create an EFA (EFA with ENA) interface. You can't use an EFA-only network interface as the primary network interface.
 - For each additional network interface, use the next unused network card index, device index 1, and either an EFA (EFA with ENA) or EFA-only network interface, depending on your use case, such as ENA bandwidth requirements or IP address space. For example use cases, see EFA configuration for a P5 instances.
 - For more information, see the EFA Guide here.

P6-B200 instances

P6-B200 contain 8 network interface cards, and can be launched using the following Amazon CLI command:

```
aws ec2 run-instances --region $REGION \
    --instance-type $INSTANCETYPE \
    --image-id $AMI --key-name $KEYNAME \
    --iam-instance-profile "Name=dlami-builder" \
    --tag-specifications "ResourceType=instance,Tags=[{Key=Name,Value=$TAG}]" \
```

```
--network-interfaces "NetworkCardIndex=0, DeviceIndex=0, Groups=$SG, SubnetId=
$SUBNET, InterfaceType=efa" \
    "NetworkCardIndex=1, DeviceIndex=1, Groups=$SG, SubnetId=$SUBNET, InterfaceType=efa" \
    "NetworkCardIndex=2, DeviceIndex=1, Groups=$SG, SubnetId=$SUBNET, InterfaceType=efa" \
    "NetworkCardIndex=3, DeviceIndex=1, Groups=$SG, SubnetId=$SUBNET, InterfaceType=efa" \
    "NetworkCardIndex=4, DeviceIndex=1, Groups=$SG, SubnetId=$SUBNET, InterfaceType=efa" \
    "NetworkCardIndex=5, DeviceIndex=1, Groups=$SG, SubnetId=$SUBNET, InterfaceType=efa" \
    "NetworkCardIndex=5, DeviceIndex=1, Groups=$SG, SubnetId=$SUBNET, InterfaceType=efa" \
    "NetworkCardIndex=6, DeviceIndex=1, Groups=$SG, SubnetId=$SUBNET, InterfaceType=efa" \
    "NetworkCardIndex=7, DeviceIndex=1, Groups=$SG, SubnetId=$SUBNET, InterfaceType=efa" \
    "NetworkCardIndex=7, DeviceIndex=1, Groups=$SG, SubnetId=$SUBNET, InterfaceType=efa" \
    "NetworkCardIndex=7, DeviceIndex=1, Groups=$SG, SubnetId=$SUBNET, InterfaceType=efa" \]
```

P5en instances

P5en contain 16 network interface cards, and can be launched using the following Amazon CLI command:

```
aws ec2 run-instances --region $REGION \
    --instance-type $INSTANCETYPE \
    --image-id $AMI --key-name $KEYNAME \
    --iam-instance-profile "Name=dlami-builder" \
    --tag-specifications "ResourceType=instance,Tags=[{Key=Name,Value=$TAG}]" \
    --network-interfaces "NetworkCardIndex=0,DeviceIndex=0,Groups=$SG,SubnetId=
$SUBNET,InterfaceType=efa" \
    "NetworkCardIndex=1,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa" \
    "NetworkCardIndex=3,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa" \
    "NetworkCardIndex=4,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa" \
    "NetworkCardIndex=15,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa" \
    "NetworkCardInde
```

P5/P5e instances

P5 and P5e instances contain 32 network interface cards, and can be launched using the following Amazon CLI command:

```
aws ec2 run-instances --region $REGION \
    --instance-type $INSTANCETYPE \
    --image-id $AMI --key-name $KEYNAME \
    --iam-instance-profile "Name=dlami-builder" \
    --tag-specifications "ResourceType=instance,Tags=[{Key=Name,Value=$TAG}]" \
    --network-interfaces "NetworkCardIndex=0,DeviceIndex=0,Groups=$SG,SubnetId=
$SUBNET,InterfaceType=efa" \
    "NetworkCardIndex=1,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa" \
```

```
"NetworkCardIndex=2,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa" \
"NetworkCardIndex=3,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa" \
"NetworkCardIndex=4,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa" \
...
"NetworkCardIndex=31,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa"
```

Kernel

• Kernel version is pinned using command:

```
echo linux-aws hold | sudo dpkg —set-selections
echo linux-headers-aws hold | sudo dpkg —set-selections
echo linux-image-aws hold | sudo dpkg —set-selections
```

 We recommend that users avoid updating their kernel version (unless due to a security patch) to ensure compatibility with installed drivers and package versions. If users still wish to update, they can run the following commands to unpin their kernel versions:

```
echo linux-aws install | sudo dpkg -set-selections
echo linux-headers-aws install | sudo dpkg -set-selections
echo linux-image-aws install | sudo dpkg -set-selections
```

• For each new version of DLAMI, latest available compatible kernel is used.

Release Date: 2025-05-16

AMI name: Deep Learning Base OSS Nvidia Driver GPU AMI (Ubuntu 22.04) 20250516

Added

Added support for P6-B200 EC2 instances

Updated

- Upgraded EFA Installer from version 1.39.0 to 1.40.0
- Upgrade Amazon OFI NCCL Plugin from version 1.13.0-aws to 1.14.2-aws
- Updated compiled NCCL Version from version 2.22.3 to 2.26.5
- Updated default CUDA version from version 12.6 to 12.8
- Updated Nvidia DCGM version from 3.3.9 to 4.4.3

Release Date: 2025-05-05

AMI name: Deep Learning Base OSS Nvidia Driver GPU AMI (Ubuntu 22.04) 20250503

Updated

• Upgraded GDRCopy from 2.4.1 to 2.5.1

Release Date: 2025-04-24

AMI name: Deep Learning Base OSS Nvidia Driver GPU AMI (Ubuntu 22.04) 20250424

Updated

 Upgraded Nvidia driver from version 570.124.06 to 570.133.20 to address CVEs present in the NVIDIA GPU Display Driver Security Bulletin for April 2025

Release Date: 2025-02-17

AMI name: Deep Learning Base OSS Nvidia Driver GPU AMI (Ubuntu 22.04) 20250214

Updated

- Updated NVIDIA Container Toolkit from version 1.17.3 to version 1.17.4
 - Please see the release notes page here for more information: <u>https://github.com/NVIDIA/</u> nvidia-container-toolkit/releases/tag/v1.17.4
 - In Container Toolkit version 1.17.4, the mounting of CUDA compat libraries is now disabled. In order to ensure compatibility with multiple CUDA versions on container workflows, please ensure you update your LD_LIBRARY_PATH to include your CUDA compatibility libraries as shown in the If you use a CUDA compatibility layer tutorial.

Removed

 Removed user space libraries cuobj and nvdisasm provided by <u>NVIDIA CUDA toolkit</u> to address CVEs present in the <u>NVIDIA CUDA Toolkit Security Bulletin for February 18, 2025</u>

Release Date: 2025-02-07

AMI name: Deep Learning Base OSS Nvidia Driver GPU AMI (Ubuntu 22.04) 20250205

Developer Guide

Added

• Added CUDA toolkit version 12.6 in directory /usr/local/cuda-12.6

Removed

- CUDA versions 12.1 and 12.2 has been removed from this DLAMI. Customers can install these versions from NVIDIA using the link below
 - https://developer.nvidia.com/cuda-toolkit-archive

Release Date: 2025-01-31

AMI name: Deep Learning Base OSS Nvidia Driver GPU AMI (Ubuntu 22.04) 20250131

Updated

- Upgraded EFA version from 1.37.0 to 1.38.0
 - EFA now bundles the Amazon OFI NCCL plugin, which can now be found in /opt/amazon/ofinccl rather than the original /opt/aws-ofi-nccl/. If updating your LD_LIBRARY_PATH variable, please ensure that you modify your OFI NCCL location properly.
- Upgraded Nvidia Container Toolkit from 1.17.3 to 1.17.4

Release Date: 2025-01-17

AMI name: Deep Learning Base OSS Nvidia Driver GPU AMI (Ubuntu 22.04) 20250117

Updated

 Upgraded Nvidia driver from version 550.127.05 to 550.144.03 to address CVEs present in the NVIDIA GPU Display Driver Security Bulletin for January 2025

Release Date: 2024-11-18

AMI name: Deep Learning Base OSS Nvidia Driver GPU AMI (Ubuntu 22.04) 20241115

Added

• Amazon FSx package for Lustre support added.

Fixed

 Due to a change in the Ubuntu kernel to address a defects in the Kernel Address Space Layout Randomization (KASLR) functionality, G4Dn/G5 instances are unable to properly initialize CUDA on the OSS Nvidia driver. In order to mitigate this issue, this DLAMI includes functionality that dynamically loads the proprietary driver for G4Dn and G5 instances. Please allow a brief initialization period for this loading in order to ensure that your instances are able to work properly.

To check the status and health of this service, you can use the following command:

```
sudo systemctl is-active dynamic_driver_load.service
active
```

Release Date: 2024-10-23

AMI name: Deep Learning Base OSS Nvidia Driver GPU AMI (Ubuntu 22.04) 20241023

Updated

 Upgraded Nvidia driver from version 550.90.07 to 550.127.05 to address CVEs present in the NVIDIA GPU Display Security Bulletin for October 2024

Release Date: 2024-10-01

AMI name: Deep Learning Base OSS Nvidia Driver GPU AMI (Ubuntu 20.04) 20240930

Updated

- Upgraded Nvidia driver and Fabric Manager from version 535.183.01 to 550.90.07
- Upgraded Nvidia Container Toolkit from version 1.16.1 to 1.16.2, addressing the security vulnerability CVE-2024-0133.
- Upgraded EFA Version from 1.32.0 to 1.34.0
- Upgraded NCCL to latest version 2.22.3 for all CUDA versions
 - CUDA 12.1, 12.2 upgraded from 2.18.5+CUDA12.2
 - CUDA 12.3 upgraded from version 2.21.5+CUDA12.4
Added

- Added CUDA toolkit version 12.4 in directory /usr/local/cuda-12.4
- Added support for **P5e EC2 instances**.

Release Date: 2024-08-19

AMI name: Deep Learning Base OSS Nvidia Driver GPU AMI (Ubuntu 22.04) 20240816

Added

• Added support for <u>G6e EC2 instance</u>.

Release Date: 2024-06-06

AMI name: Deep Learning Base OSS Nvidia Driver GPU AMI (Ubuntu 22.04) 20240606

Updated

• Updated Nvidia driver version to 535.183.01 from 535.161.08

Release Date: 2024-05-15

AMI name: Deep Learning Base OSS Nvidia Driver GPU AMI (Ubuntu 22.04) 20240513

Removed

 Amazon FSx for Lustre support has been removed in this release due to incompatibility with the latest Ubuntu 22.04 kernel versions. Support for FSx for Lustre will be reinstated once the latest kernel version is supported. Customers who require FSx for Lustre should continue to use the <u>Deep Learning Base GPU AMI (Ubuntu 20.04)</u>.

Release Date: 2024-04-29

AMI name: Deep Learning Base OSS Nvidia Driver GPU AMI (Ubuntu 22.04) 20240429

Added

• Initial release of the Deep Learning Base OSS DLAMI for Ubuntu 22.04

Amazon Deep Learning Base AMI (Amazon Linux 2)

For help getting started, see Getting started with DLAMI.

AMI name format

- Deep Learning Base OSS Nvidia Driver AMI (Amazon Linux 2) Version \${XX.X}
- Deep Learning Base Proprietary Nvidia Driver AMI (Amazon Linux 2) Version \${XX.X}

Supported EC2 instances

- Please refer to Important changes to DLAMI.
- Deep Learning with OSS Nvidia Driver supports G4dn, G5, G6, Gr6, G6e, P4d, P4de, P5, P5e, P5en
- Deep Learning with Proprietary Nvidia Driver supports G3 (G3.16x not supported), P3, P3dn

The AMI includes the following:

- Supported Amazon Service: Amazon EC2
- Operating System: Amazon Linux 2
- Compute Architecture: x86
- Latest available version is installed for the following packages:
 - Linux Kernel: 5.10
 - Docker
 - Amazon CLI v2 at /usr/local/bin/aws2 and Amazon CLI v1 at /usr/bin/aws
 - Nvidia container toolkit:
 - Version command: nvidia-container-cli -V
 - Nvidia-docker2:
 - Version command: nvidia-docker version
- **Python**: /usr/bin/python3.7
- NVIDIA Driver:
 - OSS Nvidia driver: 550.163.01
 - Proprietary Nvidia driver: 550.163.01
- NVIDIA CUDA 12.1-12.4 stack:
 - CUDA, NCCL and cuDDN installation directories: /usr/local/cuda-xx.x/

- Default CUDA: 12.1
 - PATH /usr/local/cuda points to CUDA 12.1
 - Updated below env vars:
 - LD_LIBRARY_PATH to have /usr/local/cuda-12.1/lib:/usr/local/cuda-12.1/lib64:/usr/local/cuda-12.1:/usr/local/cuda-12.1/targets/x86_64-linux/lib
 - PATH to have /usr/local/cuda-12.1/bin/:/usr/local/cuda-12.1/include/
 - For any different CUDA version, please update LD_LIBRARY_PATH accordingly.
- Compiled NCCL Version: 2.22.3
- NCCL Tests Location:
 - all_reduce, all_gather and reduce_scatter: /usr/local/cuda-xx.x/efa/test-cuda-xx.x/
 - To run NCCL tests, LD_LIBRARY_PATH needs to passed having below updates.
 - Common PATHs are already added to LD_LIBRARY_PATH:
 - /opt/amazon/efa/lib:/opt/amazon/openmpi/lib:/opt/aws-ofi-nccl/lib:/usr/local/lib:/ usr/lib
 - For any different CUDA version, please update LD_LIBRARY_PATH accordingly.
- EFA installer: 1.38.0
- Nvidia GDRCopy: 2.4
- Amazon OFI NCCL: 1.13.2
 - Amazon OFI NCCL now supports multiple NCCL versions with single build
 - Installation path: /opt/amazon/ofi-nccl/ . Path /opt/amazon/ofi-nccl/lib64 is added to LD_LIBRARY_PATH.
- EBS volume type: gp3
- Query AMI-ID with SSM Parameter (example Region is us-east-1):
 - OSS Nvidia Driver:

```
aws ssm get-parameter --region us-east-1 \
    --name /aws/service/deeplearning/ami/x86_64/base-oss-nvidia-driver-amazon-
linux-2/latest/ami-id \
    --query "Parameter.Value" \
    --output text
```

• Proprietary Nvidia Driver:

```
aws ssm get-parameter --region us-east-1 \setminus
```

```
--name /aws/service/deeplearning/ami/x86_64/base-proprietary-nvidia-driver-
amazon-linux-2/latest/ami-id \
    --query "Parameter.Value" \
    --output text
```

- Query AMI-ID with AWSCLI (example Region is us-east-1):
 - OSS Nvidia Driver:

```
aws ec2 describe-images --region us-east-1 \
    --owners amazon \
    --filters 'Name=name,Values=Deep Learning Base OSS Nvidia Driver AMI (Amazon
Linux 2) Version ??.?' 'Name=state,Values=available' \
    --query 'reverse(sort_by(Images, &CreationDate))[:1].ImageId' \
    --output text
```

• Proprietary Nvidia Driver:

```
aws ec2 describe-images --region us-east-1 \
    --owners amazon \
    --filters 'Name=name,Values=Deep Learning Base Proprietary Nvidia Driver AMI
(Amazon Linux 2) Version ??.?' 'Name=state,Values=available' \
    --query 'reverse(sort_by(Images, &CreationDate))[:1].ImageId' \
    --output text
```

Notices

NVIDIA Container Toolkit 1.17.4

In Container Toolkit version 1.17.4 the mounting of CUDA compat libraries is now disabled. In order to ensure compatibility with multiple CUDA versions on container workflows, please ensure you update your LD_LIBRARY_PATH to include your CUDA compatibility libraries as shown in the <u>If you</u> <u>use a CUDA compatibility layer</u> tutorial.

EFA Updates from 1.37 to 1.38 (Release on 2025-02-04)

EFA now bundles the Amazon OFI NCCL plugin, which can now be found in /opt/amazon/ofi-nccl rather than the original /opt/aws-ofi-nccl/. If updating your LD_LIBRARY_PATH variable, please ensure that you modify your OFI NCCL location properly.

Support policy

These AMIs Components of this AMI like CUDA versions may be removed and changed based on <u>framework support policy</u> or to optimize performance for <u>deep learning containers</u> or to reduce AMI size in a future release, without prior notice. We remove CUDA versions from AMIs if they are not used by any supported framework version.

EC2 instances with multiple network cards

- Many instances types that support EFA also have multiple network cards.
- DeviceIndex is unique to each network card, and must be a non-negative integer less than the limit of ENIs per NetworkCard. On P5, the number of ENIs per NetworkCard is 2, meaning that the only valid values for DeviceIndex is 0 or 1.
 - For the primary network interface (network card index 0, device index 0), create an EFA (EFA with ENA) interface. You can't use an EFA-only network interface as the primary network interface.
 - For each additional network interface, use the next unused network card index, device index 1, and either an EFA (EFA with ENA) or EFA-only network interface, depending on your use case, such as ENA bandwidth requirements or IP address space. For example use cases, see EFA configuration for a P5 instances.
 - For more information, see the EFA Guide here.

P5/P5e instances

• P5 and P5e instances contain 32 network interface cards, and can be launched using the following Amazon CLI command:

```
aws ec2 run-instances --region $REGION \
    --instance-type $INSTANCETYPE \
    --image-id $AMI --key-name $KEYNAME \
    --iam-instance-profile "Name=dlami-builder" \
    --tag-specifications "ResourceType=instance, Tags=[{Key=Name,Value=$TAG}]" \
    --network-interfaces "NetworkCardIndex=0,DeviceIndex=0,Groups=$SG,SubnetId=
$SUBNET,InterfaceType=efa" \
    "NetworkCardIndex=1,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa"
    "NetworkCardIndex=2,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa"
    "NetworkCardIndex=3,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa"
}
```

. . .

```
"NetworkCardIndex=4, DeviceIndex=1, Groups=$SG, SubnetId=$SUBNET, InterfaceType=efa"
```

"NetworkCardIndex=31, DeviceIndex=1, Groups=\$SG, SubnetId=\$SUBNET, InterfaceType=efa"

P5en instances

/

P5en contain 16 network interface cards, and can be launched using the following Amazon CLI command:

```
aws ec2 run-instances --region REGION \setminus
    --instance-type $INSTANCETYPE \
    --image-id $AMI --key-name $KEYNAME \
    --iam-instance-profile "Name=dlami-builder" \
    --tag-specifications "ResourceType=instance,Tags=[{Key=Name,Value=$TAG}]" \
    --network-interfaces "NetworkCardIndex=0, DeviceIndex=0, Groups=$SG, SubnetId=
$SUBNET, InterfaceType=efa" ∖
      "NetworkCardIndex=1,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa"
 /
      "NetworkCardIndex=2,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa"
 /
      "NetworkCardIndex=3,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa"
 /
      "NetworkCardIndex=4,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa"
 /
      "NetworkCardIndex=15,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa"
```

Kernel

• Kernel version is pinned using command:

```
sudo yum versionlock kernel*
```

• We recommend that users avoid updating their kernel version (unless due to a security patch) to ensure compatibility with installed drivers and package versions. If users still wish to update they can run the following commands to unpin their kernel versions:

```
sudo yum versionlock delete kernel*
sudo yum update -y
```

• For each new version of DLAMI, latest available compatible kernel is used.

Release Date: 2025-04-22

AMI names

- Deep Learning Base OSS Nvidia Driver AMI (Amazon Linux 2) Version 69.3
- Deep Learning Base Proprietary Nvidia Driver AMI (Amazon Linux 2) Version 67.0

Updated

 Upgraded Nvidia driver from version 550.144.03 to 550.163.01 to address CVEs present in the NVIDIA GPU Display Driver Security Bulletin for April 2025

Release Date: 2025-02-17

AMI names

- Deep Learning Base OSS Nvidia Driver AMI (Amazon Linux 2) Version 68.5
- Deep Learning Base Proprietary Nvidia Driver AMI (Amazon Linux 2) Version 66.3

Updated

 Updated NVIDIA Container Toolkit from version 1.17.3 to version 1.17.4. Please see the release notes page here for more information: <u>https://github.com/NVIDIA/nvidia-container-toolkit/</u><u>releases/tag/v1.17.4</u>

Removed

 Removed user space libraries cuobj and nvdisasm provided by NVIDIA CUDA toolkit to address CVEs present in the <u>NVIDIA CUDA Toolkit Security Bulletin for February 18, 2025</u>

Release Date: 2025-02-04

AMI names

• Deep Learning Base OSS Nvidia Driver AMI (Amazon Linux 2) Version 68.4

• Deep Learning Base Proprietary Nvidia Driver AMI (Amazon Linux 2) Version 66.1

Updated

• Upgraded EFA version from 1.37.0 to 1.38.0

Release Date: 2025-01-17

AMI names

- Deep Learning Base OSS Nvidia Driver AMI (Amazon Linux 2) Version 68.3
- Deep Learning Base Proprietary Nvidia Driver AMI (Amazon Linux 2) Version 66.0

Updated

 Upgraded Nvidia driver from version 550.127.05 to 550.144.03 to address CVEs present in the NVIDIA GPU Display Driver Security Bulletin for January 2025

Release Date: 2025-01-06

AMI names

- Deep Learning Base OSS Nvidia Driver AMI (Amazon Linux 2) Version 68.2
- Deep Learning Base Proprietary Nvidia Driver AMI (Amazon Linux 2) Version 65.9

Updated

- Upgraded EFA from version 1.34.0 to 1.37.0
- Upgraded Amazon OFI NCCL from version 1.11.0 to 1.13.0

Release Date: 2024-12-09

AMI names

- Deep Learning Base OSS Nvidia Driver AMI (Amazon Linux 2) Version 68.1
- Deep Learning Base Proprietary Nvidia Driver AMI (Amazon Linux 2) Version 65.8

Updated

• Upgraded Nvidia Container Toolkit from version 1.17.0 to 1.17.3

Release Date: 2024-11-09

AMI names

- Deep Learning Base OSS Nvidia Driver AMI (Amazon Linux 2) Version 67.9
- Deep Learning Base Proprietary Nvidia Driver AMI (Amazon Linux 2) Version 65.6

Updated

• Upgraded Nvidia Container Toolkit from version 1.16.2 to 1.17.0, addressing the security vulnerability CVE-2024-0134.

Release Date: 2024-10-22

AMI names

- Deep Learning Base OSS Nvidia Driver AMI (Amazon Linux 2) Version 67.7
- Deep Learning Base Proprietary Nvidia Driver AMI (Amazon Linux 2) Version 65.4

Updated

 Upgraded Nvidia driver from version 550.90.07 to 550.127.05 to address CVEs present in the NVIDIA GPU Display Security Bulletin for October 2024

Release Date: 2024-10-03

AMI names

- Deep Learning Base OSS Nvidia Driver AMI (Amazon Linux 2) Version
- Deep Learning Base Proprietary Nvidia Driver AMI (Amazon Linux 2) Version 65.2

Updated

 Upgraded Nvidia Container Toolkit from version 1.16.1 to 1.16.2, addressing the security vulnerability <u>CVE-2024-0133</u>.

Release Date: 2024-08-27

AMI name: Deep Learning Base OSS Nvidia Driver AMI (Amazon Linux 2) Version 67.0

Updated

- Upgraded Nvidia driver and Fabric Manager from version 535.183.01 to 550.90.07
 - Removed multi-user shell requirement from Fabric Manager based on Nvidia recommendations
 - Please reference known issues for Tesla driver 550.90.07 here for more information
- Upgraded EFA Version from 1.32.0 to 1.34.0
- Upgraded NCCL to latest version 2.22.3 for all CUDA versions
 - CUDA 12.1, 12.2 upgraded from 2.18.5+CUDA12.2
 - CUDA 12.3 upgraded from 2.21.5+CUDA12.4

Added

- Added CUDA toolkit version 12.4 in directory /usr/local/cuda-12.4
- Added support for P5e EC2 instances.

Removed

• Removed CUDA Toolkit version 11.8 stack present in directory /usr/local/cuda-11.8

Release Date: 2024-08-19

AMI name: Deep Learning Base OSS Nvidia Driver AMI (Amazon Linux 2) Version 66.3

Added

• Added support for G6e EC2 instances.

Release Date: 2024-06-06

AMI names

- Deep Learning Base OSS Nvidia Driver AMI (Amazon Linux 2) Version 65.4
- Deep Learning Base Proprietary Nvidia Driver AMI (Amazon Linux 2) Version 63.9

Updated

• Updated Nvidia driver version to 535.183.01 from 535.161.08

Release Date: 2024-05-02

AMI names

- Deep Learning Base OSS Nvidia Driver AMI (Amazon Linux 2) Version 64.7
- Deep Learning Base Proprietary Nvidia Driver AMI (Amazon Linux 2) Version 63.2

Updated

- Updated EFA version from version 1.30 to version 1.32
- Updated Amazon OFI NCCL plugin from version 1.7.4 to version 1.9.1
- Updated Nvidia container toolkit from version 1.13.5 to version 1.15.0

Added

• Added CUDA12.3 stack with CUDA12.3, NCCL 2.21.5, CuDNN 8.9.7

Version 1.15.0 does NOT include the nvidia-container-runtime and nvidia-docker2 packages. It is recommended to use nvidia-container-toolkit packages directly by following the <u>Nvidia container</u> toolkit docs.

Removed

• Removed CUDA11.7, CUDA12.0 stacks present at /usr/local/cuda-11.7 and /usr/local/cuda-12.0

 Removed nvidia-docker2 package and its command nvidia-docker as part of Nvidia container toolkit update from 1.13.5 to 1.15.0 which does **NOT** include the nvidia-container-runtime and nvidia-docker2 packages.

Release Date: 2024-04-04

AMI name: Deep Learning Base OSS Nvidia Driver AMI (Amazon Linux 2) Version 64.0

Added

• For OSS Nvidia driver DLAMIs, added G6 and Gr6 EC2 instances support

Release Date: 2024-03-29

AMI names

- Deep Learning Base OSS Nvidia Driver AMI (Amazon Linux 2) Version 62.3
- Deep Learning Base Proprietary Nvidia Driver AMI (Amazon Linux 2) Version 63.2

Updated

- Updated Nvidia driver from 535.104.12 to 535.161.08 in both Proprietary and OSS Nvidia driver DLAMIs.
- The new supported instances for each DLAMI are as follows:
 - Deep Learning with Proprietary Nvidia Driver supports G3 (G3.16x not supported), P3, P3dn
 - Deep Learning with OSS Nvidia Driver supports G4dn, G5, P4d, P4de, P5.

Removed

• Removed G4dn, G5, G3.16x EC2 instances support from Proprietary Nvidia driver DLAMI.

Release Date: 2024-03-20

AMI name: Deep Learning Base OSS Nvidia Driver AMI (Amazon Linux 2) Version 63.1

Added

 Added awscliv2 in the AMI as /usr/local/bin/aws2, alongside awscliv1 as /usr/local/bin/aws on OSS Nvidia Driver AMI

Release Date: 2024-03-13

AMI name: Deep Learning Base OSS Nvidia Driver AMI (Amazon Linux 2) Version 63.0

Updated

- Updated OSS Nvidia driver DLAMI with G4dn and G5 support, based on it current support looks like below:
 - Deep Learning Base Proprietary Nvidia Driver AMI (Amazon Linux 2) supports P3, P3dn, G3, G4dn, G5.
 - Deep Learning Base OSS Nvidia Driver AMI (Amazon Linux 2) supports G4dn, G5, P4, P5.
- OSS Nvidia driver DLAMIs are recommended to be used for G4dn, G5, P4, P5.

Release Date: 2024-02-13

AMI names

- Deep Learning Base OSS Nvidia Driver AMI (Amazon Linux 2) Version 62.1
- Deep Learning Base Proprietary Nvidia Driver AMI (Amazon Linux 2) Version 62.1

Updated

- Updated OSS Nvidia driver from 535.129.03 to 535.154.05
- Updated EFA from 1.29.0 to 1.30.0
- Updated Amazon OFI NCCL from 1.7.3-aws to 1.7.4-aws

Release Date: 2024-02-01

AMI name: Deep Learning Base Proprietary Nvidia Driver AMI (Amazon Linux 2) Version 62.0

Security

• Updated runc package version to consume patch for <u>CVE-2024-21626</u>.

Version 61.4

AMI name: Deep Learning Base OSS Nvidia Driver AMI (Amazon Linux 2) Version 61.4

Updated

• OSS Nvidia Driver updated from 535.104.12 to 535.129.03

Version 61.0

AMI name: Deep Learning Base OSS Nvidia Driver AMI (Amazon Linux 2) Version 61.4

Updated

- EFA updated from 1.26.1 to 1.29.0
- GDRCopy updated from 2.3 to 2.4

Added

- Amazon Deep Learning AMI (DLAMI) is split into two separate groups:
 - DLAMI that uses Nvidia Proprietary Driver (to support P3, P3dn, G3, G5, G4dn).
 - DLAMI that uses Nvidia OSS Driver to enable EFA (to support P4, P5).
- Please refer to public announcement for more information on DLAMI split.
- For Amazon CLI queries, see the bullet point Query AMI-ID with AWSCLI (example Region is useast-1)

Version 60.6

AMI name: Deep Learning Base AMI (Amazon Linux 2) Version 60.6

Updated

- Amazon OFI NCCL Plugin updated from version 1.7.2 to version 1.7.3
- Updated CUDA 12.0-12.1 directories with NCCL version 2.18.5
- CUDA12.1 updated as the default CUDA Version
 - Updated LD_LIBRARY_PATH to have /usr/local/cuda-12.1/targets/x86_64-linux/lib/:/usr/local/cuda-12.1/lib:/usr/local/cuda-12.1/lib64:/usr/local/cuda-12.1 and PATH to have /usr/local/cuda-12.1/bin/

• For customers looking to change to any different CUDA version, please define the LD_LIBRARY_PATH and PATH variables accordingly.

Added

 Kernel Live Patching is now enabled. Live patching enables customers to apply security vulnerability and critical bug patches to a running Linux kernel, without reboots or disruptions to running applications. Please note that live patching support for kernel 5.10.192 will end on 11/30/23.

Version 60.5

AMI name: Deep Learning Base AMI (Amazon Linux 2) Version 60.5

Updated

• NVIDIA Driver updated from 535.54.03 to 535.104.12

This latest driver fixes NVML ABI breaking changes found in the 535.54.03 driver, as well as the driver regression found in driver 535.86.10 that affected CUDA toolkits on P5 instances. Please reference the following NVIDIA release notes for details on the fixes:

- 4235941 NVML ABI Breaking change fix
- 4228552 CUDA Toolkit Error Fix
- Updated CUDA 12.2 directories with NCCL 2.18.5
- EFA updated from 1.24.1 to latest 1.26.1

Added

• Added CUDA12.2 at /usr/local/cuda-12.2

Removed

• Removed support for CUDA 11.5 and CUDA 11.6

Version 60.2

AMI name: Deep Learning Base AMI (Amazon Linux 2) Version 60.2

Updated

• Updated aws-ofi-nccl plugin from v1.7.1 to v1.7.2

Version 60.0

Release date: 2023-08-11

Added

- This AMI now provides support for Multi-node training functionality on P5 and all the previouslysupported EC2 instances
- For P5 EC2 instances, NCCL 2.18 is recommended to be used and has been added to CUDA12.0, and CUDA12.1.

Removed

• Removed support for CUDA11.5.

Version 59.2

Release date: 2023-08-08

Removed

• Removed CUDA-11.3 and CUDA-11.4

Version 59.1

Release date: 2023-08-03

Updated

- Updated Amazon OFI NCCL plugin to v1.7.1
- Made CUDA11.8 as default as PyTorch 2.0 supports 11.8 and for P5 EC2 instance, it is recommended to use >=CUDA11.8
 - Updated LD_LIBRARY_PATH to have /usr/local/cuda-11.8/targets/x86_64-linux/lib/:/usr/local/cuda-11.8/lib:/usr/local/cuda-11.8/lib64:/usr/local/cuda-11.8 and PATH to have /usr/local/cuda-11.8/bin/

• For any different cuda version, please define LD_LIBRARY_PATH accordingly.

Fixed

• Fixed Nvidia Fabric Manager (FM) package load issue mentioned in earlier Release Date 2023-07-19.

Version 58.9

Release date: 2023-07-19

Updated

- Updated Nvidia driver from 525.85.12 to 535.54.03
- Updated EFA installer from 1.22.1 to 1.24.1

Added

- Added c-state changes to disable idle state of processor by setting the max c-state to C1. This change is made by setting `intel_idle.max_cstate=1 processor.max_cstate=1` in the linux boot arguments in file /etc/default/grub
- Amazon EC2 P5 instance support:
 - Added P5 EC2 instance support for workflows using single node/instance. Multi-node support (e.g. for multi-node training) using EFA (Elastic Fabric Adapter) and Amazon OFI NCCL plugin will added in an upcoming release.
 - Please use CUDA>=11.8 for optimal performance.
 - Known Issue: Nvidia Fabric Manager (FM) package takes time to load on P5, customers need to wait for 2-3 minutes until FM loads after launching P5 instance. To check if FM is started, please run command sudo systemctl is-active nvidia-fabricmanager, it should return active before starting any workflow. This will be fixed in upcoming release.

Version 58.0

Release date: 2023-05-19

Removed

• Removed CUDA11.0-11.2 stack as per support policy mentioned in the top section of this document.

Version 57.3

Release date: 2023-04-06

Added

• Added Nvidia GDRCopy 2.3

Version 56.8

Release date: 2023-03-09

Updated

• Updated NVIDIA Driver from 515.65.01 to 525.85.12

Added

• Added cuda-11.8 at /usr/local/cuda-11.8/

Version 56.0

Release date: 2022-12-06

Updated

• Updated EFA version from 1.17.2 to 1.19.0

Version 55.0

Release date: 2022-11-04

Updated

• Updated NVIDIA Driver from 510.47.03 to 515.65.01

Added

• Added cuda-11.7 at /usr/local/cuda-11.7/

Version 54.0

Release date: 2022-09-15

Updated

• Updated EFA version from 1.16.0 to 1.17.2

Version 53.3

Release date: 2022-05-25

Updated

- Updated aws-efa-installer to version 1.15.2
- Updated aws-ofi-nccl to version 1.3.0-aws which include the topology for p4de.24xlarge.

Added

• This release adds support for p4de.24xlarge EC2 instances.

Version 53.0

Release date: 2022-04-28

Added

- Added Amazon CloudWatch Agent
- Added three systemd services which uses predefined json files available at path /opt/aws/ amazon-cloudwatch-agent/etc/ to configure GPU metrics using linux user cwagent
 - dlami-cloudwatch-agent@minimal
 - Commands to enable GPU metrics:

sudo systemctl enable dlami-cloudwatch-agent@minimal

sudo systemctl start dlami-cloudwatch-agent@minimal

- It creates these metrics: utilization_gpu, utilization_memory
- dlami-cloudwatch-agent@partial
 - Commands to enable GPU metrics:

```
sudo systemctl enable dlami-cloudwatch-agent@partial
sudo systemctl start dlami-cloudwatch-agent@partial
```

- It creates these metrics: utilization_gpu, utilization_memory, memory_total, memory_used, memory_free
- dlami-cloudwatch-agent@all
 - Commands to enable GPU metrics:

```
sudo systemctl enable dlami-cloudwatch-agent@all
sudo systemctl start dlami-cloudwatch-agent@all
```

• It creates all available GPU metrics

Version 52.0

Release date: 2022-03-08

Updated

Updated Kernel version to 5.10

Version 51.0

Release date: 2022-03-04

Updated

Updated Nvidia Driver to 510.47.03

Version 50.0

Release date: 2022-02-17

Updated

- Locked aws-neuron-dkms and tensorflow-model-server-neuron as they get updated to newer versions which are not supported by Neuron packages present in AMI
 - Commands if customer would like to unlock the package to update them to latest: sudo yum versionlock delete aws-neuron-dkms sudo yum versionlock delete tensorflow-model-serverneuron

Version 49.0

Release date: 2022-01-13

Added

- Added CUDA11.2 with the following components:
 - cuDNN v8.1.1.33
 - NCCL 2.8.4
 - CUDA 11.2.2

Updated

• Updated symlink pip to pip3

Deprecations

- Deprecated support for the P2 instance type
- Deprecated python2.7 and removed related python2.7 packages such as "python-dev", "pythonpip", and "python-tk"

Version 48.0

Release date: 2021-12-27

Updated

 Removed org.apache.ant_1.9.2.v201404171502\lib\ant-apache-log4j.jar from cuda versions as it is not being used and there is no risk to users who have the Log4j files. For more information, see <u>https://nvidia.custhelp.com/app/answers/detail/a_id/5294</u>.

Version 47.0

Release date: 2021-11-24

Updated

• Updated EFA to 1.14.1

Version 46.0

Release date: 2021-11-12

Updated

- Updated Neuron packages from aws-neuron-dkms=1.5.*, aws-neuron-runtime-base=1.5.*, aws-neuron-tools=1.6.* to aws-neuron-dkms=2.2.*, aws-neuron-runtime-base=1.6.*, aws-neuron-tools=2.0.*.
- Removed Neuron package aws-neuron-runtime=1.5.* as Neuron no longer have a runtime running as daemon and runtime is now integrated with framework as a library.

Version 45.0

Release date: 2021-10-21

Added

• Security scan reports in JSON format are available at /opt/aws/dlami/info/.

Version 44.0

Release date: 2021-10-08

Changed

 For every instance launch using DLAMI, tag "aws-dlami-autogenerated-tag-do-not-delete" will be added which will allow Amazon to collect instance type, instance ID, DLAMI type, and OS information. No information on the commands used within the DLAMI is collected or retained. No other information about the DLAMI is collected or retained. To opt out of usage tracking for your DLAMI, add a tag to your Amazon EC2 instance during launch. The tag should use the key

OPT_OUT_TRACKING with the associated value set to true. For more information, see <u>Tag your</u> <u>Amazon EC2 resources</u>.

Security

• Updated docker version to docker-20.10.7-3

Version 43.0

Release date: 2021-08-24

Changed

• Updated "notebook" to version "6.4.1".

Version 42.0

Release date: 2021-07-23

Changed

• Updated Nvidia driver and Fabric manager version to 450.142.00.

Version 41.0

Release date: 2021-06-24

Changed

• Updated Neuron packages as per Neuron Release v1.14.0

Version 40.0

Release date: 2021-06-10

Changed

• Updated awscli version to 1.19.89

Version 39.0

Release date: 2021-05-27

Security

• Removed vulnerable CUDA-10.0 componenets (Visual Profiler, Nsight EE, and JRE) from the CUDA-10.0 installation (/usr/local/cuda-10.0).

Version 38.0

Release date: 2021-05-25

Changed

• Upgraded runc to latest

Version 37.0

Release date: 2021-04-23

Changed

• Updated Nvidia Tesla driver and Fabric Manager version to 450.119.03.

Version 36.1

Release date: 2021-04-21

Fixed

• Fixed an issue that slowed down the instance launch speed.

Version 36.0

Release date: 2021-03-24

Added

• Added tensorflow-model-server-neuron to support neuron model serving.

Changed

• Upgraded jupyterlab to version 3.0.8 for python3.

Fixed

- The old installation of OpenMPI in /usr/local/mpi caused /opt/amazon/openmpi/bin/mpirun to be linked incorrectly. To fix the link issue, we removed /usr/local/mpi installation, OpenMPI installation in /opt/amazon/openmpi is available.
- Remove duplicated and non-existing definition of shell environments that has been polluting the shell environment variables such as PATH, and LD_LIBRARY_PATH. As the result, ~/.dlami, and / etc/profile.d/var.sh has been removed, and /etc/profile.d/dlami.sh has been added.

Security

Updated package cryptography to address <u>CVE-2020-36242</u>

Version 35.0

Release date: 2021-03-08

Added

Added TensorRT CUDA 11.0 installation

Version 34.3

Release date: 2021-02-25

Fixed

• Fixed a typo in the MOTD (message of the day) that incorrectly displayed version 34.1.

Version 34.2

Release date: 2021-02-24

Security

• Patched python2 and python3 for CVE-2021-3177

Known Issue

• There is a typo in the MOTD (message of the day) that incorrectly displayed version 34.1, we will be releasing version 34.3 to address this issue.

Version 34.0

Release date: 2021-02-09

Changed

• Pinned pip to version 20.3.4 for python2, this is the last pip version that support python2, and python3.5.

Version 33.0

Release date: 2021-01-19

Changed

• Updated cuDNN version to v8.0.5.39 in CUDA11.0 and CUDA11.1.

Version 32.0

Release date: 2020-12-01

Added

 Added CUDA11.1 with NCCL 2.7.8, cuDNN 8.0.4.30 for Deep Learning AMI (Amazon Linux 2), Deep Learning AMI (Ubuntu 16.04), Deep Learning AMI (Ubuntu 18.04), Deep Learning Base AMI (Ubuntu 16.04), Deep Learning Base AMI (Ubuntu 18.04), Deep Learning Base AMI (Amazon Linux 2).

Version 31.0

Release date: 2020-11-02

Changed

- Upgraded EFA installer to version 1.10.0.
- Upgraded cuDNN version to v8.0.4.30 for CUDA 11.0.
- Upgraded Amazon Neuron to version 1.1

Version 30.0

Release date: 2020-10-08

Changed

- Updated NVIDIA Driver and Fabric Manager versions to 450.80.02
- Updated NCCL to 2.7.8 in for CUDA11.0

Fixed

• Fixed an issue where yum managed python package overridden by pipmanaged installations. Executable pip, pip3, and pip3.7 has been moved from /usr/binto /usr/local/binas part of this fix.

Version 29.0

Release date: 2020-09-11

Changed

- Updated NVIDIA driver from version 450.51.05 to 450.51.06
- Added NVIDIA Fabric Manager version 450.51.06
- Upgraded EFA to 1.9.4

Version 28.0

Release date: 2020-08-19

Changed

• Added CUDA 11.0 stack with NCCL 2.7.6, and cuDNN 8.0.2.39

Version 27.0

Release date: 2020-08-07

Changed

- Upgraded EFA from version 1.7.1 to 1.9.3 at /opt/amazon/efa
- Upgraded Open MPI from version 4.0.3 to 4.0.4 in '/usr/local/mpi'. Open MPI at '/opt/amazon/ openmpi/bin/mpirun' is still at version 4.0.3
- Updated NVIDIA Driver from 440.33.01 to 450.51.05
- Upgraded NCCL version from 2.6.4 to 2.7.6 in CUDA10.2

Version 26.0

Release date: 2020-08-03

Changed

- Upgraded Amazon OFI NCCL to latest, see here for more detail.
- Cuda 8.0/9.0/9.2 have been removed from the AMI

Fixed

• Fixed an error where shared object file: libopencv_dnn.so.4.2 cannot be opened.

Version 25.0

Release date: 2020-07-19

Changed

- EFA version updated to 1.7.1 to support NCCL 2.6.4
- NCCL version updated to 2.6.4 for CUDA 10.2
- awscli version updated from 1.16.76 to 1.18.80
- boto3 version updated from 1.9.72 to 1.14.3

Version 24.1

Release date: 2020-06-14

Changed

• Docker version updated to 19.03.6

Version 24.0

Release date: 2020-05-20

Changed

• Docker version updated to 19.03.6

Version 23.0

Release date: 2020-04-29

Changed

• Upgraded python package versions

Version 22.0

Release date: 2020-03-04

Changed

- Added CUDA 10.2 stack
- Updated CUDA 10.0 and 10.1 for cuDNN and NCCL version

Amazon Deep Learning Base Qualcomm AMI (Amazon Linux 2)

For help getting started, see Getting started with DLAMI.

AMI name format

• Deep Learning Base Qualcomm AMI (Amazon Linux 2) \${YYYY-MM-DD}

The AMI includes the following:

- Supported Amazon Service: Amazon EC2
- Operating System: Amazon Linux 2
- Compute Architecture: x86
- Linux Kernel: 5.10.210-201.852.amzn2.x86_64
- SDK Location: /opt/qai-sdk
- QTI Utils Location: /opt/qti-aic
- Platform SDK Version: 1.12.0.88
- Apps SDK Version: 1.12.0.87
- EBS volume type: gp3
- Python: python3.8
- Supported EC2 Instances: dl2q
- Query AMI-ID with Amazon CLI (example Region is us-east-1):

```
aws ec2 describe-images --region us-east-1 \
    --owners amazon \
    --filters 'Name=name,Values=Deep Learning Base Qualcomm AMI (Amazon Linux
2) ??????' 'Name=state,Values=available' \
    --query 'reverse(sort_by(Images, &CreationDate))[:1].ImageId' \
    --output text
```

Notices

[4/12/24] Audit Package Removal

 DLAMIs released between March 26,2024 (2024-03-26) and April 12, 2024 (2024-04-12) were shipped without the audit package. If you require this specific package for your logging and monitoring needs, please migrate your workflows to the latest DLAMI in order to consume those with the audit package installed.

QAIC Environment:

• By default, the qaic-pytools are not enabled through the Qualcomm AI100 Apps SDK. In order to enable and create the qaic pip environment qaic-env please run the following commands:

```
echo 'yes' | bash /opt/qai-sdk/qaic-apps-*/uninstall.sh
cd /opt/qai-sdk/qaic-apps-*/
sudo sed -i "s/python3 -V/python3.8 -V/g" install.sh
./install.sh --enable-gaic-pytools
```

Version 20240314

Release date: 2024-03-14

AMI name: Deep Learning Base Qualcomm AMI (Amazon Linux 2) 20240314

Added

- AI 100 Platform SDK updated from version 1.10.0.200 to version 1.12.0.88
- AI 100 Apps SDK updated from version 1.10.0.193 to version 1.12.0.87
- SDK Version 1.12 adds support for transformer decoder models such as Llama-2 and Starcoder

Version 20240110

Release Date: 2024-01-10

AMI name: Deep Learning Base Qualcomm AMI (Amazon Linux 2) 20240103

Added

• Qualcomm AI 100 Platform image upgraded to 1.10.0.200

Version 20231115

Release Date: 2023-11-15

AMI name: Deep Learning Base Qualcomm AMI (Amazon Linux 2) 20231115

Added

- Initial release of Deep Learning Base Qualcomm AMI (Amazon Linux 2) series.
 - Please refer to the official Qualcomm documentation for more information on the Platform and Apps SDK: <u>https://quic.github.io/cloud-ai-sdk-pages/</u>
 - Please refer to the official Amazon documentation to learn more about dl2q instances: <u>DL2q</u> instances.

ARM64 Base DLAMI Release Notes

ARM64 Base DLAMI Release Notes

ARM64 Base DLAMI Release Notes

Below are the release notes for ARM64 Base DLAMI:

GPU

- Amazon Deep Learning Base ARM64 AMI (Amazon Linux 2023) (Supports P6e-GB200)
- Amazon Deep Learning Base ARM64 AMI (Ubuntu 22.04) (Supports P6e-GB200)
- Amazon Deep Learning Base ARM64 AMI (Amazon Linux 2)

Amazon Neuron

• Refer to the <u>Neuron DLAMI User Guide</u>.

Amazon Deep Learning ARM64 Base GPU AMI (Amazon Linux 2023)

For help getting started, see Getting started with DLAMI.

AMI name format

• Deep Learning ARM64 Base OSS Nvidia Driver GPU AMI (Amazon Linux 2023) \${YYYY-MM-DD}

Supported EC2 instances

• G5g, P6e-GB200(CUDA>=12.8 is supported on P6e-GB200)

The AMI includes the following:

- Supported Amazon Service: Amazon EC2
- Operating System: Amazon Linux 2023
- Compute Architecture: ARM64
- Latest available version is installed for the following packages:
 - Linux Kernel: 6.12

- FSx Lustre
- Docker
- Amazon CLI v2 at /usr/bin/aws
- NVIDIA DCGM
- Nvidia container toolkit:
 - Version command: nvidia-container-cli -V
- Nvidia-docker2:
 - Version command: nvidia-docker version
- NVIDIA Driver: 570.158.01
- NVIDIA CUDA 12.4, 12.5, 12.6, 12.8 stack:
 - CUDA, NCCL and cuDDN installation directories: /usr/local/cuda-xx.x/
 - Example: /usr/local/cuda-12.8/ , /usr/local/cuda-12.8/
 - Compiled NCCL Version:
 - For CUDA directory of 12.4, compiled NCCL Version 2.22.3+CUDA12.4
 - For CUDA directory of 12.5, compiled NCCL Version 2.22.3+CUDA12.5
 - For CUDA directory of 12.6, compiled NCCL Version 2.24.3+CUDA12.6
 - For CUDA directory of 12.8, compiled NCCL Version 2.27.5+CUDA12.8
 - Default CUDA: 12.8
 - PATH /usr/local/cuda points to CUDA 12.8
 - Updated below env vars:
 - LD_LIBRARY_PATH to have /usr/local/cuda-12.8/lib:/usr/local/cuda-12.8/lib64:/usr/local/cuda-12.8/targets/sbsa-linux/lib:/usr/local/cuda-12.8/nvvm/lib64:/usr/local/cuda-12.8/extras/CUPTI/lib64
 - PATH to have /usr/local/cuda-12.8/bin/:/usr/local/cuda-12.8/include/
 - For any different CUDA version, please update LD_LIBRARY_PATH accordingly.
- EFA installer: 1.42.0
- Nvidia GDRCopy: 2.5.1
- Amazon OFI NCCL plugin comes with EFA installer
 - Paths /opt/amazon/ofi-nccl/lib and /opt/amazon/ofi-nccl/efa are added to LD_LIBRARY_PATH.
- Amazon CLI v2 at /usr/local/bin/aws

- EBS volume type: gp3
- Python: /usr/bin/python3.9
- Query AMI-ID with SSM Parameter (example region is us-east-1):

```
SSM_PARAMETER=base-oss-nvidia-driver-gpu-amazon-linux-2023/latest/ami-id \
    aws ssm get-parameter --region us-east-1 \
    --name /aws/service/deeplearning/ami/arm64/$SSM_PARAMETER \
    --query "Parameter.Value" \
    --output text
```

• Query AMI-ID with AWSCLI (example region is us-east-1):

```
aws ec2 describe-images --region us-east-1 --owners amazon --filters
'Name=name,Values=Deep Learning ARM64 Base OSS Nvidia Driver GPU AMI (Amazon Linux
2023) ???????' 'Name=state,Values=available' --query 'reverse(sort_by(Images,
&CreationDate))[:1].ImageId' --output text
```

P6e-GB200 instances

P6e-GB200 instances contain 17 network interface cards, and can be launched using the following Amazon CLI command:

```
aws ec2 run-instances --region $REGION \
    --instance-type $INSTANCETYPE \
    --image-id $AMI --key-name $KEYNAME \
    --iam-instance-profile "Name=dlami-builder" \
    --tag-specifications "ResourceType=instance,Tags=[{Key=Name,Value=$TAG}]" \
    --network-interfaces \
        "NetworkCardIndex=0,DeviceIndex=0,Groups=$SG,SubnetId=
$SUBNET, InterfaceType=interface" \
        "NetworkCardIndex=1, DeviceIndex=0, Groups=$SG, SubnetId=
$SUBNET, InterfaceType=efa-only" ∖
        "NetworkCardIndex=2, DeviceIndex=0, Groups=$SG, SubnetId=
$SUBNET, InterfaceType=efa-only" \
        "NetworkCardIndex=3, DeviceIndex=0, Groups=$SG, SubnetId=
$SUBNET,InterfaceType=efa-only" \
        "NetworkCardIndex=4, DeviceIndex=0, Groups=$SG, SubnetId=
$SUBNET,InterfaceType=efa-only" \
        "NetworkCardIndex=5, DeviceIndex=0, Groups=$SG, SubnetId=
$SUBNET,InterfaceType=efa-only" \
```

"NetworkCardIndex=6,DeviceIndex=0,Groups=\$SG,SubnetId=
<pre>\$SUBNET,InterfaceType=efa-only" \</pre>
"NetworkCardIndex=7,DeviceIndex=0,Groups=\$SG,SubnetId=
<pre>\$SUBNET,InterfaceType=efa-only" \</pre>
"NetworkCardIndex=8,DeviceIndex=0,Groups=\$SG,SubnetId=
<pre>\$SUBNET,InterfaceType=efa-only" \</pre>
"NetworkCardIndex=9,DeviceIndex=0,Groups=\$SG,SubnetId=
<pre>\$SUBNET,InterfaceType=efa-only" \</pre>
"NetworkCardIndex=10,DeviceIndex=0,Groups=\$SG,SubnetId=
<pre>\$SUBNET,InterfaceType=efa-only" \</pre>
"NetworkCardIndex=11,DeviceIndex=0,Groups=\$SG,SubnetId=
<pre>\$SUBNET,InterfaceType=efa-only" \</pre>
"NetworkCardIndex=12,DeviceIndex=0,Groups=\$SG,SubnetId=
<pre>\$SUBNET,InterfaceType=efa-only" \</pre>
"NetworkCardIndex=13,DeviceIndex=0,Groups=\$SG,SubnetId=
<pre>\$SUBNET,InterfaceType=efa-only" \</pre>
"NetworkCardIndex=14,DeviceIndex=0,Groups=\$SG,SubnetId=
<pre>\$SUBNET,InterfaceType=efa-only" \</pre>
"NetworkCardIndex=15,DeviceIndex=0,Groups=\$SG,SubnetId=
<pre>\$SUBNET,InterfaceType=efa-only" \</pre>
"NetworkCardIndex=16,DeviceIndex=0,Groups=\$SG,SubnetId=
<pre>\$SUBNET,InterfaceType=efa-only"</pre>

Notices

NVIDIA Container Toolkit 1.17.4

In Container Toolkit version 1.17.4 the mounting of CUDA compat libraries is now disabled. In order to ensure compatibility with multiple CUDA versions on container workflows, please ensure you update your LD_LIBRARY_PATH to include your CUDA compatibility libraries as shown in the <u>If you</u> use a CUDA compatibility layer tutorial.

Support policy

These AMIs Components of this AMI like CUDA versions may be removed and changed based on <u>framework support policy</u> or to optimize performance for <u>deep learning containers</u> or to reduce AMI size in a future release , without prior notice. We remove CUDA versions from AMIs if they are not used by any supported framework version.

Kernel

• Kernel version is pinned using command:

sudo dnf versionlock kernel*

• We recommend that users avoid updating their kernel version (unless due to a security patch) to ensure compatibility with installed drivers and package versions. If users still wish to update they can run the following commands to unpin their kernel versions:

```
sudo dnf versionlock delete kernel*
sudo dnf update -y
```

• For each new version of DLAMI, latest available compatible kernel is used.

Release Date: 2025-07-04

AMI name: Deep Learning ARM64 Base OSS Nvidia Driver GPU AMI (Amazon Linux 2023) 20250704

Updated

- Added support to EC2 instance P6e-GB200. Please note that CUDA>=12.8 is supported on P6e-GB200
- Add EFA 1.42.0
- Upgraded Nvidia driver from version 570.133.20 to 570.158.01
- Upgraded CUDA 12.8 stack with NCCL 2.27.5

Release Date: 2025-04-24

AMI name: Deep Learning ARM64 Base OSS Nvidia Driver GPU AMI (Amazon Linux 2023) 20250424

Updated

- Upgraded Nvidia driver from version 570.86.15 to 570.133.20 to address CVEs present in the NVIDIA GPU Display Driver Security Bulletin for April 2025
- Updated CUDA12.8 stack with NCCL 2.26.2
- Updated default CUDA from 12.6 to 12.8
Release Date: 2025-04-22

AMI name: Deep Learning ARM64 Base OSS Nvidia Driver GPU AMI (Amazon Linux 2023) 20250421

Updated

 Upgraded Nvidia driver from version 570.124.06 to 570.133.20 to address CVEs present in the NVIDIA GPU Display Driver Security Bulletin for April 2025

Release Date: 2025-04-04

AMI name: Deep Learning ARM64 Base OSS Nvidia Driver GPU AMI (Amazon Linux 2023) 20250404

Updated

• Kernel version updated from 6.1 to 6.12

Release Date: 2025-03-03

AMI name: Deep Learning ARM64 Base OSS Nvidia Driver GPU AMI (Amazon Linux 2023) 20250303

Updated

- Nvidia driver from 550.144.03 to 570.86.15
- Default CUDA is changed from CUDA12.4 to CUDA12.6

Added

- CUDA directory of 12.5 with compiled NCCL Version 2.22.3+CUDA12.5 and CuDNN 9.7.1.26
- CUDA directory of 12.6 with compiled NCCL Version 2.24.3+CUDA12.6 and CuDNN 9.7.1.26
- CUDA directory of 12.8 with compiled NCCL Version 2.25.1+CUDA12.8 and CuDNN 9.7.1.26

Release Date: 2025-02-14

AMI name: Deep Learning ARM64 Base OSS Nvidia Driver GPU AMI (Amazon Linux 2023) 20250214

Added

• Initial release of the Deep Learning ARM64 Base OSS DLAMI for Amazon Linux 2023

Amazon Deep Learning ARM64 Base GPU AMI (Ubuntu 22.04)

For help getting started, see <u>Getting started with DLAMI</u>.

AMI name format

• Deep Learning ARM64 Base OSS Nvidia Driver GPU AMI (Ubuntu 22.04) \${YYYY-MM-DD}

Supported EC2 instances

• G5g, P6e-GB200 (CUDA>=12.8 is supported on P6e-GB200)

The AMI includes the following:

- Supported Amazon Service: Amazon EC2
- Operating System: Ubuntu 22.04
- Compute Architecture: ARM64
- Latest available version is installed for the following packages:
 - Linux Kernel: 6.8
 - FSx Lustre
 - Docker
 - Amazon CLI v2 at /usr/bin/aws
 - NVIDIA DCGM
 - Nvidia container toolkit:
 - Version command: nvidia-container-cli -V
 - Nvidia-docker2:
 - Version command: nvidia-docker version
- NVIDIA Driver: 570.158.01
- NVIDIA CUDA 12.4, 12.5, 12.6, 12.8 stack:

• CUDA, NCCL and cuDDN installation directories: /usr/local/cuda-xx.x/

- Example: /usr/local/cuda-12.8/ , /usr/local/cuda-12.8/
- Compiled NCCL Version:
 - For CUDA directory of 12.4, compiled NCCL Version 2.22.3+CUDA12.4
 - For CUDA directory of 12.5, compiled NCCL Version 2.22.3+CUDA12.5
 - For CUDA directory of 12.6, compiled NCCL Version 2.24.3+CUDA12.6
 - For CUDA directory of 12.8, compiled NCCL Version 2.27.5+CUDA12.8
- Default CUDA: 12.8
 - PATH /usr/local/cuda points to CUDA 12.8
 - Updated below env vars:
 - LD_LIBRARY_PATH to have /usr/local/cuda-12.8/lib:/usr/local/cuda-12.8/lib64:/usr/local/cuda-12.8/targets/sbsa-linux/lib:/usr/local/cuda-12.8/nvvm/lib64:/usr/local/cuda-12.8/extras/CUPTI/lib64
 - PATH to have /usr/local/cuda-12.8/bin/:/usr/local/cuda-12.8/include/
 - For any different CUDA version, please update LD_LIBRARY_PATH accordingly.
- EFA installer: 1.42.0
- Nvidia GDRCopy: 2.5.1
- Amazon OFI NCCL plugin comes with EFA installer
 - Paths /opt/amazon/ofi-nccl/lib and /opt/amazon/ofi-nccl/efa are added to LD_LIBRARY_PATH.
- Amazon CLI v2 at /usr/local/bin/aws2 and Amazon CLI v1 at /usr/bin/aws
- EBS volume type: gp3
- Python: /usr/bin/python3.10
- Query AMI-ID with SSM Parameter (example Region is us-east-1):

```
SSM_PARAMETER=base-oss-nvidia-driver-gpu-ubuntu-22.04/latest/ami-id \
    aws ssm get-parameter --region us-east-1 \
    --name /aws/service/deeplearning/ami/arm64/$SSM_PARAMETER \
    --query "Parameter.Value" \
    --output text
```

• Query AMI-ID with AWSCLI (example Region is us-east-1):

```
--owners amazon --filters 'Name=name,Values=Deep Learning ARM64 Base OSS Nvidia
Driver GPU AMI (Ubuntu 22.04) ???????' 'Name=state,Values=available' \
    --query 'reverse(sort_by(Images, &CreationDate))[:1].ImageId' \
    --output text
```

P6e-GB200 instances

P6e-GB200 instances contain 17 network interface cards, and can be launched using the following Amazon CLI command:

```
aws ec2 run-instances --region REGION \
    --instance-type $INSTANCETYPE \
    --image-id $AMI --key-name $KEYNAME \
    --iam-instance-profile "Name=dlami-builder" \
    --tag-specifications "ResourceType=instance,Tags=[{Key=Name,Value=$TAG}]" \
    --network-interfaces \
        "NetworkCardIndex=0, DeviceIndex=0, Groups=$SG, SubnetId=
$SUBNET,InterfaceType=interface" \
        "NetworkCardIndex=1, DeviceIndex=0, Groups=$SG, SubnetId=
$SUBNET, InterfaceType=efa-only" \
        "NetworkCardIndex=2, DeviceIndex=0, Groups=$SG, SubnetId=
$SUBNET, InterfaceType=efa-only" \
        "NetworkCardIndex=3, DeviceIndex=0, Groups=$SG, SubnetId=
$SUBNET, InterfaceType=efa-only" \
        "NetworkCardIndex=4, DeviceIndex=0, Groups=$SG, SubnetId=
$SUBNET, InterfaceType=efa-only" \
        "NetworkCardIndex=5, DeviceIndex=0, Groups=$SG, SubnetId=
$SUBNET,InterfaceType=efa-only" \
        "NetworkCardIndex=6, DeviceIndex=0, Groups=$SG, SubnetId=
$SUBNET,InterfaceType=efa-only" \
        "NetworkCardIndex=7, DeviceIndex=0, Groups=$SG, SubnetId=
$SUBNET,InterfaceType=efa-only" \
        "NetworkCardIndex=8, DeviceIndex=0, Groups=$SG, SubnetId=
$SUBNET,InterfaceType=efa-only" \
        "NetworkCardIndex=9, DeviceIndex=0, Groups=$SG, SubnetId=
$SUBNET,InterfaceType=efa-only" \
        "NetworkCardIndex=10, DeviceIndex=0, Groups=$SG, SubnetId=
$SUBNET,InterfaceType=efa-only" \
        "NetworkCardIndex=11, DeviceIndex=0, Groups=$SG, SubnetId=
$SUBNET,InterfaceType=efa-only" \
        "NetworkCardIndex=12, DeviceIndex=0, Groups=$SG, SubnetId=
$SUBNET,InterfaceType=efa-only" \
```

Notices

NVIDIA Container Toolkit 1.17.4

In Container Toolkit version 1.17.4 the mounting of CUDA compat libraries is now disabled. In order to ensure compatibility with multiple CUDA versions on container workflows, please ensure you update your LD_LIBRARY_PATH to include your CUDA compatibility libraries as shown in the <u>If you</u> <u>use a CUDA compatibility layer</u> tutorial.

Multi ENI support

- Ubuntu 22.04 automatically sets up and configures source routing on multiple NICs via cloud-init on its initial boot. If your workflow includes attaching/detaching your ENI's while an instance is stopped, an additional configuration must be added to the cloud-init user data to ensure proper configuration of the NIC's during these events. A sample of the cloud config is provided below.
- Please reference this Canonical documentation here for more information on how to configure the cloud config for your instances - <u>https://documentation.ubuntu.com/aws/en/latest/aws-how-to/instances/automatically-setup-multiple-nics/</u>

```
#cloud-config
# apply network config on every boot and hotplug event
updates:
    network:
    when: ['boot', 'hotplug']
```

Support policy

These AMIs Components of this AMI like CUDA versions may be removed and changed based on <u>framework support policy</u> or to optimize performance for <u>deep learning containers</u> or to reduce AMI size in a future release , without prior notice. We remove CUDA versions from AMIs if they are not used by any supported framework version.

Kernel

• Kernel version is pinned using command:

```
echo linux-aws hold | sudo dpkg —set-selections
echo linux-headers-aws hold | sudo dpkg —set-selections
echo linux-image-aws hold | sudo dpkg —set-selections
```

• We recommend that users avoid updating their kernel version (unless due to a security patch) to ensure compatibility with installed drivers and package versions. If users still wish to update they can run the following commands to unpin their kernel versions:

```
echo linux-aws install | sudo dpkg -set-selections
echo linux-headers-aws install | sudo dpkg -set-selections
echo linux-image-aws install | sudo dpkg -set-selections
```

• For each new version of DLAMI, latest available compatible kernel is used.

Release Date: 2025-07-04

AMI name: Deep Learning ARM64 Base OSS Nvidia Driver GPU AMI (Ubuntu 22.04) 20250704

Updated

- Added support to EC2 instance P6e-GB200. Please note that CUDA>=12.8 is supported on P6e-GB200
- Add EFA 1.42.0
- Upgraded Nvidia driver from version 570.133.20 to 570.158.01
- Upgraded CUDA 12.8 stack with NCCL 2.27.5

Release Date: 2025-04-24

AMI name: Deep Learning ARM64 Base OSS Nvidia Driver GPU AMI (Ubuntu 22.04) 20250424

Updated

- Upgraded Nvidia driver from version 570.86.15 to 570.133.20 to address CVE's present in the NVIDIA GPU Display Driver Security Bulletin for April 2025
- Updated CUDA 12.8 stack with NCCL 2.26.2

- Updated default CUDA from 12.6 to 12.8
- Removed CUDA 12.3

Release Date: 2025-03-03

AMI name: Deep Learning ARM64 Base OSS Nvidia Driver GPU AMI (Ubuntu 22.04) 20250303

Updated

- Nvidia driver from 550.144.03 to 570.86.15
- Default CUDA is changed from CUDA12.1 to CUDA12.6

Added

- CUDA directory of 12.4 with compiled NCCL Version 2.22.3+CUDA12.4 and CuDNN 9.7.1.26
- CUDA directory of 12.5 with compiled NCCL Version 2.22.3+CUDA12.5 and CuDNN 9.7.1.26
- CUDA directory of 12.6 with compiled NCCL Version 2.24.3+CUDA12.6 and CuDNN 9.7.1.26
- CUDA directory of 12.8 with compiled NCCL Version 2.25.1+CUDA12.8 and CuDNN 9.7.1.26

Removed

• CUDA directory of 12.1 and 12.2

Release Date: 2025-02-17

AMI name: Deep Learning ARM64 Base OSS Nvidia Driver GPU AMI (Ubuntu 22.04) 20250214

Updated

- Updated NVIDIA Container Toolkit from version 1.17.3 to version 1.17.4
 - Please see the release notes page here for more information: <u>https://github.com/NVIDIA/</u> nvidia-container-toolkit/releases/tag/v1.17.4
 - In Container Toolkit version 1.17.4, the mounting of CUDA compat libraries is now disabled. In order to ensure compatibility with multiple CUDA versions on container workflows, please ensure you update your LD_LIBRARY_PATH to include your CUDA compatibility libraries as shown in the If you use a CUDA compatibility layer tutorial.

Removed

 Removed user space libraries cuobj and nvdisasm provided by <u>NVIDIA CUDA toolkit</u> to address CVEs present in the NVIDIA CUDA Toolkit Security Bulletin for February 18, 2025

Release Date: 2025-01-17

AMI name: Deep Learning ARM64 Base OSS Nvidia Driver GPU AMI (Ubuntu 22.04) 20250117

Updated

• Upgraded Nvidia driver from version 550.127.05 to 550.144.03 to address CVEs present in the NVIDIA GPU Display Driver Security Bulletin for January 2025

Release Date: 2024-10-23

AMI name: Deep Learning ARM64 Base OSS Nvidia Driver GPU AMI (Ubuntu 22.04) 20241023

Updated

 Upgraded Nvidia driver from version 550.90.07 to 550.127.05 to address CVEs present in the NVIDIA GPU Display Security Bulletin for October 2024

Release Date: 2024-06-06

AMI name: Deep Learning ARM64 Base OSS Nvidia Driver GPU AMI (Ubuntu 22.04) 20240606

Updated

• Updated Nvidia driver version to 535.183.01 from 535.161.08

Release Date: 2024-05-15

AMI name: Deep Learning ARM64 Base OSS Nvidia Driver GPU AMI (Ubuntu 22.04) 20240514

Added

• Initial release of the Deep Learning ARM64 Base OSS DLAMI for Ubuntu 22.04

Amazon Deep Learning ARM64 Base GPU AMI (Amazon Linux 2)

For help getting started, see Getting started with DLAMI.

AMI name format

• Deep Learning ARM64 Base OSS Nvidia Driver GPU AMI (Amazon Linux 2) \${YYYY-MM-DD}

Supported EC2 instances

• G5g

The AMI includes the following:

- Supported Amazon Service: Amazon EC2
- Operating System: Amazon Linux 2
- Compute Architecture: ARM64
- Linux Kernel: 5.10
- NVIDIA Driver: 550.144.03
- NVIDIA CUDA12.1, 12.2, 12.3 stack:
 - CUDA, NCCL and cuDDN installation directories:
 - Example: /usr/local/cuda-12.1/ , /usr/local/cuda-12.1/
 - Compiled NCCL Version:
 - For CUDA directory of 12.3, compiled NCCL Version 2.21.5+CUDA12.4
 - For CUDA directory of 12.1, 12.2, compiled NCCL Version .18.5+CUDA12.2
 - Default CUDA: 12.1
 - PATH /usr/local/cuda points to CUDA 12.1
 - Updated below env vars:
 - LD_LIBRARY_PATH to have /usr/local/cuda-12.1/lib:/usr/local/cuda-12.1/lib64:/usr/local/cuda-12.1/targets/sbsa-linux/lib:/usr/local/cuda-12.1/nvvm/lib64:/usr/local/cuda-12.1/extras/CUPTI/lib64
 - PATH to have /usr/local/cuda-12.1/bin/:/usr/local/cuda-12.1/include/
 - For any different CUDA version, please update LD_LIBRARY_PATH accordingly.
- Amazon CLI v2 at /usr/local/bin/aws2 and Amazon CLI v1 at /usr/bin/aws

- EBS volume type: gp3
- Nvidia container toolkit: 1.16.2
 - Version command: nvidia-container-cli -V
- Docker: 26.1.2
- Python: /usr/bin/python3.10
- Query AMI-ID with SSM Parameter (example Region is us-east-1):

```
SSM_PARAMETER=base-oss-nvidia-driver-gpu-amazon-linux-2/latest/ami-id \
    aws ssm get-parameter --region us-east-1 \
    --name /aws/service/deeplearning/ami/arm64/$SSM_PARAMETER \
    --query "Parameter.Value" \
    --output text
```

• Query AMI-ID with AWSCLI (example Region is us-east-1):

```
aws ec2 describe-images --region us-east-1 \
    -owners amazon \
    --filters 'Name=name,Values=Deep Learning ARM64 Base OSS Nvidia Driver GPU AMI
 (Amazon Linux 2) ???????' 'Name=state,Values=available' \
    --query 'reverse(sort_by(Images, &CreationDate))[:1].ImageId' \
    --output text
```

Notices

NVIDIA Container Toolkit 1.17.4

In Container Toolkit version 1.17.4 the mounting of CUDA compat libraries is now disabled. In order to ensure compatibility with multiple CUDA versions on container workflows, please ensure you update your LD_LIBRARY_PATH to include your CUDA compatibility libraries as shown in the <u>If you</u> <u>use a CUDA compatibility layer</u> tutorial.

Support policy

These AMIs Components of this AMI like CUDA versions may be removed and changed based on <u>framework support policy</u> or to optimize performance for <u>deep learning containers</u> or to reduce AMI size in a future release , without prior notice. We remove CUDA versions from AMIs if they are not used by any supported framework version.

Kernel

• Kernel version is pinned using command:

sudo yum versionlock kernel*

• We recommend that users avoid updating their kernel version (unless due to a security patch) to ensure compatibility with installed drivers and package versions. If users still wish to update they can run the following commands to unpin their kernel versions:

```
sudo yum versionlock delete kernel*
sudo yum update -y
```

• For each new version of DLAMI, latest available compatible kernel is used.

Release Date: 2025-02-17

AMI name: Deep Learning ARM64 Base OSS Nvidia Driver GPU AMI (Amazon Linux 2) 20250214

Updated

- Updated NVIDIA Container Toolkit from version 1.17.3 to version 1.17.4
 - Please see the release notes page here for more information: <u>https://github.com/NVIDIA/</u> nvidia-container-toolkit/releases/tag/v1.17.4
 - In Container Toolkit version 1.17.4, the mounting of CUDA compat libraries is now disabled. In order to ensure compatibility with multiple CUDA versions on container workflows, please ensure you update your LD_LIBRARY_PATH to include your CUDA compatibility libraries as shown in the <u>If you use a CUDA compatibility layer</u> tutorial.

Removed

 Removed user space libraries cuobj and nvdisasm provided by <u>NVIDIA CUDA toolkit</u> to address CVEs present in the <u>NVIDIA CUDA Toolkit Security Bulletin for February 18, 2025</u>

Release Date: 2025-01-17

AMI name: Deep Learning ARM64 Base OSS Nvidia Driver GPU AMI (Amazon Linux 2) 20250117

Updated

• Upgraded Nvidia driver from version 550.127.05 to 550.144.03 to address CVEs present in the NVIDIA GPU Display Driver Security Bulletin for January 2025

Release Date: 2024-10-22

AMI name: Deep Learning ARM64 Base OSS Nvidia Driver GPU AMI (Amazon Linux 2) 20241022

Updated

 Upgraded Nvidia driver from version 550.90.07 to 550.127.05 to address CVEs present in the NVIDIA GPU Display Security Bulletin for October 2024

Release Date: 2024-10-08

AMI name: Deep Learning ARM64 Base OSS Nvidia Driver GPU AMI (Amazon Linux 2) 20241008

Updated

• Upgraded Nvidia Container Toolkit from version 1.16.1 to 1.16.2, addressing the security vulnerability CVE-2024-0133.

Release Date: 2024-06-06

AMI name: Deep Learning ARM64 Base OSS Nvidia Driver GPU AMI (Amazon Linux 2) 20240606

Updated

• Updated Nvidia driver version to 535.183.01 from 535.161.08

Release Date: 2024-05-14

AMI name: Deep Learning ARM64 Base OSS Nvidia Driver GPU AMI (Amazon Linux 2) 20240514

Added

• Initial release of the Deep Learning ARM64 Base OSS DLAMI for Amazon Linux 2

Release Notes for Single Framework DLAMIs

Single Framework DLAMI Release Notes

- PyTorch DLAMIs
- <u>TensorFlow DLAMIs</u>

PyTorch DLAMIs

Multi Framework DLAMI Release Notes

- X86 PyTorch DLAMI Release Notes
- ARM64 PyTorch DLAMI Release Notes

X86 PyTorch DLAMI Release Notes

Below Are the Release notes for X86 PyTorch DLAMIs:

GPU

- Amazon Deep Learning AMI GPU PyTorch 2.7 (Amazon Linux 2023)
- Amazon Deep Learning AMI GPU PyTorch 2.7 (Ubuntu 22.04)
- Amazon Deep Learning AMI GPU PyTorch 2.6 (Amazon Linux 2023)
- Amazon Deep Learning AMI GPU PyTorch 2.6 (Ubuntu 22.04)
- Amazon Deep Learning AMI GPU PyTorch 2.5 (Amazon Linux 2023)
- Amazon Deep Learning AMI GPU PyTorch 2.5 (Ubuntu 22.04)
- Amazon Deep Learning AMI GPU PyTorch 2.4 (Ubuntu 22.04)

Amazon Neuron

• Refer to the Neuron DLAMI User Guide

Amazon Deep Learning OSS AMI GPU PyTorch 2.7 (Amazon Linux 2023)

For help getting started, see Getting started with DLAMI.

AMI name format

• Deep Learning OSS Nvidia Driver AMI GPU PyTorch 2.7 (Amazon Linux 2023) \${YYYY-MM-DD}

Supported EC2 instances

- Please refer to Important changes to DLAMI
- G4dn, G5, G5, Gr6, P4, P4de, P5, P5e, P5en, P6-B200

The AMI includes the following:

- Supported Amazon Service: Amazon EC2
- Operating System: Amazon Linux 2023
- Compute Architecture: x86
- Linux Kernel: 6.1
- NVIDIA Driver: 570.133.20
- NVIDIA CUDA 12.8 stack:
 - CUDA, NCCL and cuDDN installation directories: /usr/local/cuda-12.8/
 - NCCL Tests Location:
 - all_reduce, all_gather, and reduce_scatter:

/usr/local/cuda-12.8/efa/test-cuda-12.8/

- To run NCCL tests, LD_LIBRARY_PATH is already updated with needed paths.
 - Common PATHs are already added to LD_LIBRARY_PATH:

```
/opt/amazon/efa/lib:/opt/amazon/openmpi/lib:/opt/amazon/ofi-nccl/lib:/usr/
local/lib:/usr/lib
```

• LD_LIBRARY_PAT is updated with CUDA version paths:

```
/usr/local/cuda/lib:/usr/local/cuda/lib64:/usr/local/cuda:/usr/local/cuda/
targets/x86_64-linux/lib
```

Compiled NCCL Version:

```
    For CUDA directory of 12.8, compiled NCCL Version 2.26.2+CUDA12.8

PyTorch DLAMIs
```

- Default CUDA: 12.8
 - PATH /usr/local/cuda points to CUDA 12.8
 - Updated below env vars:
 - LD_LIBRARY_PATH to have /usr/local/cuda/lib:/usr/local/cuda/lib64:/usr/local/cuda/ targets/x86_64-linux/lib
 - PATH to have /usr/local/cuda/bin/:/usr/local/cuda/include/
- EFA Installer: 1.40.0
- Nvidia GDRCopy: 2.5
- Amazon OFI NCCL: 1.14.2-aws
 - Installation path: /opt/amazon/ofi-nccl/. Path /opt/amazon/ofi-nccl/lib is added to LD_LIBRARY_PATH
- Amazon CLI v2 at /usr/local/bin/aws
- EBS volume type: gp3
- Nvidia container toolkit: 1.17.7
 - Version command: nvidia-container-cli -V
- Docker: 25.0.8
- Python: /usr/bin/python3.12
- Query AMI-ID with SSM Parameter (example region is us-east-1):

```
aws ssm get-parameter --region us-east-1 \
    --name /aws/service/deeplearning/ami/x86_64/oss-nvidia-driver-gpu-pytorch-2.7-
amazon-linux-2023/latest/ami-id \
    --query "Parameter.Value" \
    --output text
```

• Query AMI-ID with AWSCLI (example region is us-east-1):

```
aws ec2 describe-images --region us-east-1 --owners amazon --filters
'Name=name,Values=Deep Learning OSS Nvidia Driver AMI GPU PyTorch 2.7 (Amazon Linux
2023) ???????' 'Name=state,Values=available' --query 'reverse(sort_by(Images,
&CreationDate))[:1].ImageId' --output text
```

Notices

P6-B200 Instances

- P6-B200 instances require CUDA version 12.8 or higher and NVIDIA driver 570 or newer drivers.
- P6-B200 contain 8 network interface cards and can be launched using the following Amazon CLI command:

```
aws ec2 run-instances --region $REGION \
    --instance-type $INSTANCETYPE \
    --image-id $AMI --key-name $KEYNAME \
    --iam-instance-profile "Name=dlami-builder" \
    --tag-specifications "ResourceType=instanace,Tags=[{Key=Name,Value=$TAG}]" \
    --network-interfaces ""NetworkCardIndex=0,DeviceIndex=0,Groups=$SG,SubnetId=
$SUBNET,InterfaceType=efa" \
    "NetworkCardIndex=1,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa" \
    "NetworkCardIndex=3,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa" \
    "NetworkCardIndex=4,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa" \
    "NetworkCardIndex=4,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa" \
    "NetworkCardIndex=4,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa" \
    "NetworkCardIndex=7,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa" \\
    "NetworkCardIndex=7,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa" \\
    "NetworkCardIndex=7,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=Efa" \\
    "NetworkCardIndex=7,D
```

P5/P5e Instances

 DeviceIndex is unique to each NetworkCard and must be a non-negative integer less than the limit of ENIs per NetworkCard. On P5, the number of ENIs per NetworkCard is 2, meaning that the only valid values for DeviceIndex are 0 or 1. Below is an example of EC2 P5 instance launch command using awscli showing NetworkCardIndex for numbers 0-31 and DeviceIndex as 0 for the first interface and 1 for the remaining 31 interfaces.

```
aws ec2 run-instances --region $REGION \
    --instance-type $INSTANCETYPE \
    --image-id $AMI --key-name $KEYNAME \
    --iam-instance-profile "Name=dlami-builder" \
    --tag-specifications "ResourceType=instanace,Tags=[{Key=Name,Value=$TAG}]" \
    --network-interfaces ""NetworkCardIndex=0,DeviceIndex=0,Groups=$SG,SubnetId=
$SUBNET,InterfaceType=efa" \
    "NetworkCardIndex=1,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa" \
    "NetworkCardIndex=3,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa" \
    "NetworkCardIndex=3,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa" \
    "NetworkCardIndex=4,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa" \
    "NetworkCardIndex=4,Devi
```

```
....
```

```
"NetworkCardIndex=31, DeviceIndex=1, Groups=$SG, SubnetId=$SUBNET, InterfaceType=efa"
```

Kernel

• Kernel version is pinned using command:

```
sudo dnf versionlock kernel*
```

• We recommend users to avoid updating their kernel version (unless due to a security patch) to ensure compatibility with installed drivers and package versions. If users still wish to update they can run the following commands to unpin their kernel versions:

```
sudo dnf versionlock delete kernel*
sudo dnf update -y
```

• For each new version of DLAMI, latest available compatible kernel is used.

PyTorch Deprecation of Anaconda Channel

Starting with PyTorch 2.6, PyTorch has deprecated support for Conda (see <u>official announcement</u>). As a result, PyTorch 2.6 and above will move to using Python Virtual Environments. To activate the PyTorch venv, please use source /opt/pytorch/bin/activate

Release Date: 2025-05-22

AMI name: Deep Learning OSS Nvidia Driver AMI GPU PyTorch 2.7 (Amazon Linux 2023) 20250520

Added

 Initial release of Deep Learning AMI GPU PyTorch 2.7 (Amazon Linux 2023) series. Including a Python virtual environment pytorch (source /opt/pytorch/bin/activate) complimented with NVIDIA Driver R570, CUDA=12.8, cuDNN=9.10, PyTorch NCCL=2.26.2, and EFA=1.40.0.

Amazon Deep Learning OSS AMI GPU PyTorch 2.7 (Ubuntu 22.04)

For help getting started, see Getting started with DLAMI.

AMI name format

• Deep Learning OSS Nvidia Driver AMI GPU PyTorch 2.7 (Ubuntu 22.04) \${YYYY-MM-DD}

Supported EC2 instances

- Please refer to Important changes to DLAMI
- G4dn, G5, G5, Gr6, P4, P4de, P5, P5e, P5en, P6-B200

The AMI includes the following:

- Supported Amazon Service: Amazon EC2
- Operating System: Ubuntu 22.04
- Compute Architecture: x86
- Linux Kernel: 6.8
- NVIDIA Driver: 570.133.20
- NVIDIA CUDA 12.8 stack:
 - CUDA, NCCL and cuDDN installation directories: /usr/local/cuda-12.8/
 - NCCL Tests Location:
 - all_reduce, all_gather, and reduce_scatter:

/usr/local/cuda-12.8/efa/test-cuda-12.8/

- To run NCCL tests, LD_LIBRARY_PATH is already updated with needed paths.
 - Common PATHs are already added to LD_LIBRARY_PATH:

```
/opt/amazon/efa/lib:/opt/amazon/openmpi/lib:/opt/amazon/ofi-nccl/lib:/usr/
local/lib:/usr/lib
```

• LD_LIBRARY_PAT is updated with CUDA version paths:

```
/usr/local/cuda/lib:/usr/local/cuda/lib64:/usr/local/cuda:/usr/local/cuda/
targets/x86_64-linux/lib
```

Compiled NCCL Version:

```
    For CUDA directory of 12.8, compiled NCCL Version 2.26.2+CUDA12.8

PyTorch DLAMIs
```

- Default CUDA: 12.8
 - PATH /usr/local/cuda points to CUDA 12.8
 - Updated below env vars:
 - LD_LIBRARY_PATH to have /usr/local/cuda/lib:/usr/local/cuda/lib64:/usr/local/cuda/ targets/x86_64-linux/lib
 - PATH to have /usr/local/cuda/bin/:/usr/local/cuda/include/
- EFA Installer: 1.40.0
- Nvidia GDRCopy: 2.5
- Nvidia Transformer Engine: 1.11.0
- Amazon OFI NCCL: 1.14.2-aws
 - Installation path: /opt/amazon/ofi-nccl/. Path /opt/amazon/ofi-nccl/lib is added to LD_LIBRARY_PATH
- Amazon CLI v2 at /usr/local/bin/aws
- EBS volume type: gp3
- Nvidia container toolkit: 1.17.7
 - Version command: nvidia-container-cli -V
- Docker: 28.2.2
- **Python**: /usr/bin/python3.12
- Query AMI-ID with SSM Parameter (example region is us-east-1):

```
aws ssm get-parameter --region us-east-1 \
    --name /aws/service/deeplearning/ami/x86_64/oss-nvidia-driver-gpu-pytorch-2.7-
ubuntu-22.04/latest/ami-id \
    --query "Parameter.Value" \
    --output text
```

• Query AMI-ID with AWSCLI (example region is us-east-1):

```
aws ec2 describe-images --region us-east-1 --owners amazon --filters
'Name=name,Values=Deep Learning OSS Nvidia Driver AMI GPU PyTorch 2.7 (Ubuntu
22.04) ??????' 'Name=state,Values=available' --query 'reverse(sort_by(Images,
&CreationDate))[:1].ImageId' --output text
```

Notices

Flash Attention

- Flash attention does not yet have an <u>official release for PyTorch 2.7</u>. For this reason, it is temporarily removed from this AMI. Once an official release is made for PyTorch 2.7, we will include it in this AMI.
- Without flash attention, transformer engine defaults to using cuDNN fused attention. There are currently known issues with fused attention and Blackwell GPUs, like P6-B200 instances.
 - "With compute capability sm10.0 (Blackwell-architecture) GPUs, the FP8 datatype with scaled dot-product attention contains a deadlock that causes the kernel to hang under some circumstances, such as when the problem size is large or the GPU is running multiple kernels simultaneously. A fix is planned for a future release." [cuDNN 9.10.0 release notes]
 - For users seeking to run P6-B200 instances with FP8 data and scaled dot-product attention, please consider installing flash attention manually.

P6-B200 Instances

- P6-B200 instances require CUDA version 12.8 or higher and NVIDIA driver 570 or newer drivers.
- P6-B200 contain 8 network interface cards and can be launched using the following Amazon CLI command:

```
aws ec2 run-instances --region $REGION \
    --instance-type $INSTANCETYPE \
    --image-id $AMI --key-name $KEYNAME \
    --iam-instance-profile "Name=dlami-builder" \
    --tag-specifications "ResourceType=instanace,Tags=[{Key=Name,Value=$TAG}]" \
    --network-interfaces ""NetworkCardIndex=0,DeviceIndex=0,Groups=$SG,SubnetId=
$SUBNET,InterfaceType=efa" \
    "NetworkCardIndex=1,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa" \
    "NetworkCardIndex=3,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa" \
    "NetworkCardIndex=3,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa" \
    "NetworkCardIndex=4,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa" \
    ....
    "NetworkCardIndex=7,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa" \
    "NetworkCardInd
```

P5/P5e Instances

 DeviceIndex is unique to each NetworkCard and must be a non-negative integer less than the limit of ENIs per NetworkCard. On P5, the number of ENIs per NetworkCard is 2, meaning that the only valid values for DeviceIndex are 0 or 1. Below is an example of EC2 P5 instance launch command using awscli showing NetworkCardIndex for numbers 0-31 and DeviceIndex as 0 for the first interface and 1 for the remaining 31 interfaces.

```
aws ec2 run-instances --region $REGION \
    --instance-type $INSTANCETYPE \
    --image-id $AMI --key-name $KEYNAME \
    --iam-instance-profile "Name=dlami-builder" \
    --tag-specifications "ResourceType=instanace,Tags=[{Key=Name,Value=$TAG}]" \
    --network-interfaces ""NetworkCardIndex=0,DeviceIndex=0,Groups=$SG,SubnetId=
$SUBNET,InterfaceType=efa" \
    "NetworkCardIndex=1,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa" \
    "NetworkCardIndex=2,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa" \
    "NetworkCardIndex=3,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa" \
    "NetworkCardIndex=4,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa" \
    "NetworkCardIndex=4,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa" \
    "NetworkCardIndex=3,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa" \
    "NetworkCardIndex=3,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa" \
    "NetworkCardIndex=3,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa" \
    "NetworkCardIndex=3,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa" \
    "NetworkCardIndex=3,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa" \
    "NetworkCardIndex=31,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa" \
    "NetworkCardIndex=31,DeviceIndex=1,Groups=SSG,SubnetId=$SUBNET,InterfaceType=Efa" \
    "NetworkCardInd
```

Kernel

• Kernel version is pinned using command:

```
echo linux-aws hold | sudo dkpg -set-selections
echo linux-headers-aws hold | sudo dpkg -set-selections
echo linux-image-aws hold | sudo dpkg -set-selections
```

• We recommend users to avoid updating their kernel version (unless due to a security patch) to ensure compatibility with installed drivers and package versions. If users still wish to update they can run the following commands to unpin their kernel versions:

```
echo linux-aws install | sudo dpkg -set-selections
echo linux-headers-aws install | sudo dpkg -set-selections
echo linux-image-aws install | sudo dpkg -set-selections
apt-get upgrade -y
```

• For each new version of DLAMI, latest available compatible kernel is used.

PyTorch Deprecation of Anaconda Channel

Starting with PyTorch 2.6, PyTorch has deprecated support for Conda (see <u>official announcement</u>). As a result, PyTorch 2.6 and above will move to using Python Virtual Environments. To activate the PyTorch venv, please use source /opt/pytorch/bin/activate

Release Date: 2025-06-03

AMI name: Deep Learning OSS Nvidia Driver AMI GPU PyTorch 2.7 (Ubuntu 22.04) 20250602

Added

Initial release of Deep Learning AMI GPU PyTorch 2.7 (Ubuntu 22.04) series. Including a Python virtual environment pytorch (source /opt/pytorch/bin/activate) complimented with NVIDIA Driver R570, CUDA=12.8, cuDNN=9.10, PyTorch NCCL=2.26.5, and EFA=1.40.0.

Known Issues

- "With compute capability sm10.0 (Blackwell-architecture) GPUs, the FP8 datatype with scaled dot-product attention contains a deadlock that causes the kernel to hang under some circumstances, such as when the problem size is large or the GPU is running multiple kernels simultaneously. A fix is planned for a future release." [cuDNN 9.10.0 release notes]
 - For users seeking to run P6-B200 instances with FP8 data and scaled dot-product attention, please consider installing flash attention manually.

Amazon Deep Learning AMI GPU PyTorch 2.6 (Amazon Linux 2023)

For help getting started, see <u>Getting started with DLAMI</u>.

AMI name format

• Deep Learning OSS NVIDIA Driver AMI GPU PyTorch 2.6.0 (Amazon Linux 2023) \${YYYY-MM-DD}

Supported EC2 instances:

Please refer to Important changes to DLAMI

 Deep Learning with OSS NVIDIA Driver supports G4dn, G5, G6, Gr6, G6e, P4d, P4de, P5, P5e, P5en

The AMI includes the following:

- Supported Amazon Service: EC2
- Operating System: Amazon Linux 2023
- Compute Architecture: x86
- NVIDIA CUDA12.6 stack:
 - CUDA, NCCL and cuDDN installation path: /usr/local/cuda-12.6/
 - Default CUDA: 12.6
 - PATH /usr/local/cuda points to /usr/local/cuda-12.6/
 - Updated below env vars:
 - LD_LIBRARY_PATH to have /usr/local/cuda/lib:/usr/local/cuda/lib64:/usr/local/cuda:/usr/ local/cud/targets/x86_64-linux/lib
 - PATH to have /usr/local/cuda/bin/:/usr/local/cuda/include/
 - Compiled NCCL Version for 12.6: 2.24.3
- NCCL Tests Location:
 - all_reduce, all_gather and reduce_scatter: /usr/local/cuda-xx.x/efa/test-cuda-xx.x/
 - To run NCCL tests, LD_LIBRARY_PATH is already with updated with needed paths.
 - Common PATHs are already added to LD_LIBRARY_PATH:
 - /opt/amazon/efa/lib:/opt/amazon/openmpi/lib:/opt/aws-ofi-nccl/lib:/usr/local/lib:/usr/ lib
 - LD_LIBRARY_PATH is updated with CUDA version paths
 - /usr/local/cuda/lib:/usr/local/cuda/lib64:/usr/local/cuda:/usr/local/cud/targets/x86_64linux/lib
- EFA Installer: 1.38.0
- Nvidia GDRCopy: 2.4.1
- Amazon OFI NCCL: 1.13.2-aws
 - Amazon OFI NCCL now supports multiple NCCL versions with single build
 - Installation path: /opt/amazon/ofi-nccl/ . Path /opt/amazon/ofi-nccl/lib is added to LD_LIBRARY_PATH.

- Python version: 3.12
- Python: /opt/pytorch/bin/python
- NVIDIA Driver: 570.86.15
- Amazon CLI v2 at /usr/bin/aws
- EBS volume type: gp3
- NVMe Instance Store Location (on Supported EC2 instances): /opt/dlami/nvme
- Query AMI-ID with SSM Parameter (example Region is us-east-1):
 - OSS Nvidia Driver:

```
aws ssm get-parameter --region us-east-1 \
    --name /aws/service/deeplearning/ami/x86_64/oss-nvidia-driver-gpu-pytorch-2.6-
amazon-linux-2023/latest/ami-id \
    --query "Parameter.Value" \
    --output text
```

- Query AMI-ID with AWSCLI (example Region is us-east-1):
 - OSS Nvidia Driver:

```
aws ec2 describe-images --region us-east-1 \
    --owners amazon --filters 'Name=name,Values=Deep Learning OSS Nvidia Driver AMI
GPU PyTorch 2.6.? (Amazon Linux 2023) ???????' 'Name=state,Values=available' \
    --query 'reverse(sort_by(Images, &CreationDate))[:1].ImageId' \
    --output text
```

Notices

PyTorch Deprecation of Anaconda Channel

Starting with PyTorch 2.6, PyTorch has deprecated support for Conda (see <u>official announcement</u>). As a result, PyTorch 2.6 and above will move to using Python Virtual Environments. To activate the PyTorch venv, please use source /opt/pytorch/bin/activate **P5/P5e Instances:**

 DeviceIndex is unique to each NetworkCard, and must be a non-negative integer less than the limit of ENIs per NetworkCard. On P5, the number of ENIs per NetworkCard is 2, meaning that the only valid values for DeviceIndex is 0 or 1. Below is the example of EC2 P5 instance launch command using awscli showing NetworkCardIndex from number 0-31 and DeviceIndex as 0 for first interface and DeviceIndex as 1 for rest 31 interrfaces.

```
aws ec2 run-instances --region $REGION \
    --instance-type $INSTANCETYPE \
    --image-id $AMI --key-name $KEYNAME \
    --iam-instance-profile "Name=dlami-builder" \
    --tag-specifications "ResourceType=instance,Tags=[{Key=Name,Value=$TAG}]" \
    --network-interfaces "NetworkCardIndex=0, DeviceIndex=0, Groups=$SG, SubnetId=
$SUBNET,InterfaceType=efa" \
      "NetworkCardIndex=1,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa"
\
      "NetworkCardIndex=2, DeviceIndex=1, Groups=$SG, SubnetId=$SUBNET, InterfaceType=efa"
 /
      "NetworkCardIndex=3,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa"
/
      "NetworkCardIndex=4,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa"
 /
      . . .
      "NetworkCardIndex=31, DeviceIndex=1, Groups=$SG, SubnetId=$SUBNET, InterfaceType=efa"
```

Kernel

• Kernel version is pinned using command:

```
sudo dnf versionlock kernel*
```

• We recommend users to avoid updating their kernel version (unless due to a security patch) to ensure compatibility with installed drivers and package versions. If users still wish to update they can run the following commands to unpin their kernel versions:

```
sudo dnf versionlock delete kernel*
sudo dnf update -y
```

• For each new version of DLAMI, latest available compatible kernel is used.

Release Date: 2025-02-21

AMI name: Deep Learning OSS Nvidia Driver AMI GPU PyTorch 2.6.0 (Amazon Linux 2023) 20250220

Added

- Initial release of the Deep Learning OSS Nvidia Driver AMI GPU PyTorch 2.6 for Amazon Linux 2023
 - As of PyTorch2.6, Pytorch has deprecated support for Conda. As a result, Pytorch 2.6 and above will move to using Python Virtual Environments. To activate the pytorch venv, please use source /opt/pytorch/bin/activate

Amazon Deep Learning AMI GPU PyTorch 2.6 (Ubuntu 22.04)

For help getting started, see Getting started with DLAMI.

AMI name format

 Deep Learning OSS Nvidia Driver AMI GPU PyTorch 2.6.\${PATCH-VERSION} (Ubuntu 22.04) \${YYYY-MM-DD}

Supported EC2 instances

- Please refer to Important changes to DLAMI.
- Deep Learning with OSS Nvidia Driver supports G4dn, G5, G6, Gr6, P4, P4de, P5, P5e, P5en.

The AMI includes the following:

- Supported Amazon Service: Amazon EC2
- Operating System: Ubuntu 22.04
- Compute Architecture: x86
- **Python**: /opt/pytorch/bin/python
- NVIDIA Driver:
 - OSS Nvidia driver: 570.86.15
- NVIDIA CUDA12.1 stack:
 - CUDA, NCCL and cuDDN installation path: /usr/local/cuda-12.6/
 - Default CUDA: 12.6
 - PATH /usr/local/cuda points to /usr/local/cuda-12.6/
 - Updated below env vars:

- LD_LIBRARY_PATH to have /usr/local/cuda/lib:/usr/local/cuda/lib64:/usr/local/cuda:/usr/ local/cuda/targets/x86_64-linux/lib
- PATH to have /usr/local/cuda/bin/:/usr/local/cuda/include/
- Compiled system NCCL Version present at /usr/local/cuda/: 2.24.3
- PyTorch Compiled NCCL Version from PyTorch conda environment: 2.21.5
- NCCL Tests Location:
 - all_reduce, all_gather and reduce_scatter: /usr/local/cuda-xx.x/efa/test-cuda-xx.x/
 - To run NCCL tests, LD_LIBRARY_PATH is already with updated with needed paths.
 - Common PATHs are already added to LD_LIBRARY_PATH:
 - /opt/amazon/efa/lib:/opt/amazon/openmpi/lib:/opt/aws-ofi-nccl/lib:/usr/local/lib:/usr/lib
 - LD_LIBRARY_PATH is updated with CUDA version paths
 - /usr/local/cuda/lib:/usr/local/cuda/lib64:/usr/local/cuda:/usr/local/cud/targets/x86_64linux/lib
- EFA Installer: 1.38.0
- Nvidia GDRCopy: 2.4.1
- Nvidia Transformer Engine: v1.11.0
- Amazon OFI NCCL: 1.13.2-aws
 - Installation path: /opt/aws-ofi-nccl/ . Path /opt/aws-ofi-nccl/lib is added to LD_LIBRARY_PATH.
 - Note: PyTorch package comes with dynamically linked Amazon OFI NCCL plugin as a conda package aws-ofi-nccl-dlc package as well and PyTorch will use that package instead of system Amazon OFI NCCL.
- Amazon CLI v2 as aws2 and Amazon CLI v1 as aws
- EBS volume type: gp3
- Python version: 3.11
- Query AMI-ID with SSM Parameter (example Region is us-east-1):
 - OSS Nvidia Driver:

--output text

- Query AMI-ID with AWSCLI (example Region is us-east-1):
 - OSS Nvidia Driver:

```
aws ec2 describe-images --region us-east-1 \
    --owners amazon --filters 'Name=name,Values=Deep Learning OSS Nvidia Driver AMI
GPU PyTorch 2.6.? (Ubuntu 22.04) ???????' 'Name=state,Values=available' \
    --query 'reverse(sort_by(Images, &CreationDate))[:1].ImageId' \
    --output text
```

Notices

PyTorch Deprecation of Anaconda Channel

Starting with PyTorch2.6, Pytorch has deprecated support for Conda (see <u>official announcement</u>). As a result, Pytorch 2.6 and above will move to using Python Virtual Environments. Tto activate the pytorch venv, please use source /opt/pytorch/bin/activate

P5/P5e Instances:

 DeviceIndex is unique to each NetworkCard, and must be a non-negative integer less than the limit of ENIs per NetworkCard. On P5, the number of ENIs per NetworkCard is 2, meaning that the only valid values for DeviceIndex is 0 or 1. Below is the example of EC2 P5 instance launch command using awscli showing NetworkCardIndex from number 0-31 and DeviceIndex as 0 for first interface and DeviceIndex as 1 for rest 31 interrfaces.

```
aws ec2 run-instances --region $REGION \
    --instance-type $INSTANCETYPE \
    --image-id $AMI --key-name $KEYNAME \
    --iam-instance-profile "Name=dlami-builder" \
    --tag-specifications "ResourceType=instance,Tags=[{Key=Name,Value=$TAG}]" \
    --network-interfaces "NetworkCardIndex=0,DeviceIndex=0,Groups=$SG,SubnetId=
$SUBNET,InterfaceType=efa" \
    "NetworkCardIndex=1,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa"
    "NetworkCardIndex=2,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa"
    "NetworkCardIndex=3,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa"
}
```

. . .

```
"NetworkCardIndex=4,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa"
```

```
"NetworkCardIndex=31, DeviceIndex=1, Groups=$SG, SubnetId=$SUBNET, InterfaceType=efa"
```

Kernel

/

• Kernel version is pinned using command:

```
echo linux-aws hold | sudo dpkg —set-selections
echo linux-headers-aws hold | sudo dpkg —set-selections
echo linux-image-aws hold | sudo dpkg —set-selections
```

• We recommend users to avoid updating their kernel version (unless due to a security patch) to ensure compatibility with installed drivers and package versions. If users still wish to update they can run the following commands to unpin their kernel versions:

```
echo linux-aws install | sudo dpkg -set-selections
echo linux-headers-aws install | sudo dpkg -set-selections
echo linux-image-aws install | sudo dpkg -set-selections
apt-get upgrade -y
```

• For each new version of DLAMI, latest available compatible kernel is used.

Release Date: 2025-02-21

AMI name: Deep Learning OSS Nvidia Driver AMI GPU PyTorch 2.6.0 (Ubuntu 22.04) 20250220

Added

- Initial release of Deep Learning AMI GPU PyTorch 2.6 (Ubuntu 22.04) series. Including a Python virtual environment pytorch (source /opt/pytorch/bin/activate), complimented with NVIDIA Driver R570, CUDA=12.6, cuDNN=9.7, PyTorch NCCL=2.21.5, and EFA=1.38.0.
 - Starting with PyTorch2.6, Pytorch has deprecated support for Conda (see <u>official</u> <u>announcement</u>). As a result, Pytorch 2.6 and above will move to using Python Virtual Environments. Tto activate the pytorch venv, please activate using source /opt/pytorch/bin/ activate

Amazon Deep Learning AMI GPU PyTorch 2.5 (Amazon Linux 2023)

For help getting started, see Getting started with DLAMI.

AMI name format

• Deep Learning OSS Nvidia Driver AMI GPU PyTorch 2.5.1 (Amazon Linux 2023) \${YYYY-MM-DD}

Supported EC2 instances

- Please refer to Important changes to DLAMI.
- Deep Learning with OSS Nvidia Driver supports G4dn, G5, G6, Gr6, G6e, P4d, P4de, P5, P5e, P5en

The AMI includes the following:

- Supported Amazon Service: EC2
- Operating System: Amazon Linux 2023
- Compute Architecture: x86
- NVIDIA CUDA12.4 stack:
 - CUDA, NCCL and cuDDN installation path: /usr/local/cuda-12.4/
 - Default CUDA: 12.4
 - PATH /usr/local/cuda points to /usr/local/cuda-12.4/
 - Updated below env vars:
 - LD_LIBRARY_PATH to have /usr/local/cuda/lib:/usr/local/cuda/lib64:/usr/local/cuda:/usr/ local/cud/targets/x86_64-linux/lib
 - PATH to have /usr/local/cuda/bin/:/usr/local/cuda/include/
 - Compiled NCCL Version for 12.4: 2.21.5
- NCCL Tests Location:
 - all_reduce, all_gather and reduce_scatter: /usr/local/cuda-xx.x/efa/test-cuda-xx.x/
 - To run NCCL tests, LD_LIBRARY_PATH is already with updated with needed paths.
 - Common PATHs are already added to LD_LIBRARY_PATH:
 - /opt/amazon/efa/lib:/opt/amazon/openmpi/lib:/opt/aws-ofi-nccl/lib:/usr/local/lib:/usr/ lib
 - LD_LIBRARY_PATH is updated with CUDA version paths

- /usr/local/cuda/lib:/usr/local/cuda/lib64:/usr/local/cuda:/usr/local/cud/targets/x86_64linux/lib
- EFA Installer: 1.38.0
- Nvidia GDRCopy: 2.4.1
- Amazon OFI NCCL: 1.13.2-aws
 - Amazon OFI NCCL now supports multiple NCCL versions with single build
 - Installation path: /opt/aws-ofi-nccl/ . Path /opt/aws-ofi-nccl/lib is added to LD_LIBRARY_PATH.
 - Tests path for ring, message_transfer: /opt/aws-ofi-nccl/tests
- Python version: 3.11
- Python: /opt/conda/envs/pytorch/bin/python
- NVIDIA Driver: 560.35.03
- Amazon CLI v2 at /usr/bin/aws
- EBS volume type: gp3
- NVMe Instance Store Location (on Supported EC2 Instances): /opt/dlami/nvme
- Query AMI-ID with SSM Parameter (example Region is us-east-1):
 - OSS Nvidia Driver:

- Query AMI-ID with AWSCLI (example Region is us-east-1):
 - OSS Nvidia Driver:

```
aws ec2 describe-images --region us-east-1 \
    --owners amazon --filters 'Name=name,Values=Deep Learning OSS Nvidia Driver AMI
GPU PyTorch 2.5.? (Amazon Linux 2023) ???????' 'Name=state,Values=available' \
    --query 'reverse(sort_by(Images, &CreationDate))[:1].ImageId' \
    --output text
```

Notices

P5/P5e Instances:

 DeviceIndex is unique to each NetworkCard, and must be a non-negative integer less than the limit of ENIs per NetworkCard. On P5, the number of ENIs per NetworkCard is 2, meaning that the only valid values for DeviceIndex is 0 or 1. Below is the example of EC2 P5 instance launch command using awscli showing NetworkCardIndex from number 0-31 and DeviceIndex as 0 for first interface and DeviceIndex as 1 for rest 31 interrfaces.

```
aws ec2 run-instances --region $REGION \
    --instance-type $INSTANCETYPE \
    --image-id $AMI --key-name $KEYNAME \
    --iam-instance-profile "Name=dlami-builder" \
    --tag-specifications "ResourceType=instance,Tags=[{Key=Name,Value=$TAG}]" \
    --network-interfaces "NetworkCardIndex=0, DeviceIndex=0, Groups=$SG, SubnetId=
$SUBNET, InterfaceType=efa" \
      "NetworkCardIndex=1,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa"
 /
      "NetworkCardIndex=2,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa"
 /
      "NetworkCardIndex=3,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa"
 Ι
      "NetworkCardIndex=4,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa"
 \
      "NetworkCardIndex=31, DeviceIndex=1, Groups=$SG, SubnetId=$SUBNET, InterfaceType=efa"
```

Kernel

• Kernel version is pinned using command:

```
sudo dnf versionlock kernel*
```

• We recommend users to avoid updating their kernel version (unless due to a security patch) to ensure compatibility with installed drivers and package versions. If users still wish to update they can run the following commands to unpin their kernel versions:

```
sudo dnf versionlock delete kernel*
sudo dnf update -y
```

• For each new version of DLAMI, latest available compatible kernel is used.

Release Date: 2025-02-17

AMI name: Deep Learning OSS Nvidia Driver AMI GPU PyTorch 2.5.1 (Amazon Linux 2023) 20250216

Updated

- Updated NVIDIA Container Toolkit from version 1.17.3 to version 1.17.4
 - Please see the release notes page here for more information: <u>https://github.com/NVIDIA/</u> nvidia-container-toolkit/releases/tag/v1.17.4
 - In Container Toolkit version 1.17.4, the mounting of CUDA compat libraries is now disabled. In order to ensure compatibility with multiple CUDA versions on container workflows, please ensure you update your LD_LIBRARY_PATH to include your CUDA compatibility libraries as shown in the <u>If you use a CUDA compatibility layer</u> tutorial.

Removed

 Removed user space libraries cuobj and nvdisasm provided by <u>NVIDIA CUDA toolkit</u> to address CVEs present in the <u>NVIDIA CUDA Toolkit Security Bulletin for February 18, 2025</u>

Release Date: 2025-01-08

AMI name: Deep Learning OSS Nvidia Driver AMI GPU PyTorch 2.5.1 (Amazon Linux 2023) 20250107

Added

• Added Support for <u>G4dn instances</u>.

Release Date: 2024-11-21

AMI name: Deep Learning OSS Nvidia Driver AMI GPU PyTorch 2.5.1 (Amazon Linux 2023) 20241120

Added

 Initial release of the Deep Learning OSS Nvidia Driver AMI GPU PyTorch 2.5 for Amazon Linux 2023

Known Issues

This DLAMI does not support G4dn and G5 EC2 instances at this time. Amazon is aware of an incompatibility that may result in CUDA initialization failures, affecting both G4dn and G5 instance families when using the open source NVIDIA drivers together with a Linux kernel version 6.1 or newer. This issue affects Linux distributions such as Amazon Linux 2023, Ubuntu 22.04 or newer, or SUSE Linux Enterprise Server 15 SP6 or newer, among others.

Amazon Deep Learning AMI GPU PyTorch 2.5 (Ubuntu 22.04)

For help getting started, see Getting started with DLAMI.

AMI name format

 Deep Learning OSS Nvidia Driver AMI GPU PyTorch 2.5.\${PATCH_VERSION} (Ubuntu 22.04) \${YYYY-MM-DD}

Supported EC2 instances

- Please refer to Important changes to DLAMI.
- Deep Learning with OSS Nvidia Driver supports G4dn, G5, G6, Gr6, P4, P4de, P5, P5e, P5en.

The AMI includes the following:

- Supported Amazon Service: Amazon EC2
- Operating System: Ubuntu 22.04
- **Compute Architecture**: x86
- Python: /opt/conda/envs/pytorch/bin/python
- NVIDIA Driver:
 - OSS Nvidia driver: 550.144.03
- NVIDIA CUDA12.4 stack:

- CUDA, NCCL and cuDDN installation path: /usr/local/cuda-12.4/
- Default CUDA: 12.4
 - PATH /usr/local/cuda points to /usr/local/cuda-12.4/
 - Updated below env vars:
 - LD_LIBRARY_PATH to have /usr/local/cuda/lib:/usr/local/cuda/lib64:/usr/local/cuda:/usr/ local/cuda/targets/x86_64-linux/lib
 - PATH to have /usr/local/cuda/bin/:/usr/local/cuda/include/
- Compiled system NCCL Version present at /usr/local/cuda/: 2.21.5
- PyTorch Compiled NCCL Version from PyTorch conda environment: 2.21.5
- NCCL Tests Location:
 - all_reduce, all_gather and reduce_scatter: /usr/local/cuda-xx.x/efa/test-cuda-xx.x/
 - To run NCCL tests, LD_LIBRARY_PATH is already with updated with needed paths.
 - Common PATHs are already added to LD_LIBRARY_PATH:
 - /opt/amazon/efa/lib:/opt/amazon/openmpi/lib:/opt/aws-ofi-nccl/lib:/usr/local/lib:/usr/ lib
 - LD_LIBRARY_PATH is updated with CUDA version paths
 - /usr/local/cuda/lib:/usr/local/cuda/lib64:/usr/local/cuda:/usr/local/cud/targets/x86_64linux/lib
- EFA Installer: 1.34.0
- Nvidia GDRCopy: 2.4.1
- Nvidia Transformer Engine: v1.11.0
- Amazon OFI NCCL: 1.11.0-aws
 - Installation path: /opt/aws-ofi-nccl/ . Path /opt/aws-ofi-nccl/lib is added to LD_LIBRARY_PATH.
 - Tests path for ring, message_transfer: /opt/aws-ofi-nccl/tests
 - Note: PyTorch package comes with dynamically linked Amazon OFI NCCL plugin as a conda package aws-ofi-nccl-dlc package as well and PyTorch will use that package instead of system Amazon OFI NCCL.
- Amazon CLI v2 as aws2 and Amazon CLI v1 as aws
- EBS volume type: gp3

PyTBythonsversion: 3.11

- Query AMI-ID with SSM Parameter (example Region is us-east-1):
 - OSS Nvidia Driver:

- Query AMI-ID with AWSCLI (example Region is us-east-1):
 - OSS Nvidia Driver:

```
aws ec2 describe-images --region us-east-1 \
    --owners amazon --filters 'Name=name,Values=Deep Learning OSS Nvidia Driver AMI
GPU PyTorch 2.5.? (Ubuntu 22.04) ??????' 'Name=state,Values=available' \
    --query 'reverse(sort_by(Images, &CreationDate))[:1].ImageId' \
    --output text
```

Notices

P5/P5e Instances:

 DeviceIndex is unique to each NetworkCard, and must be a non-negative integer less than the limit of ENIs per NetworkCard. On P5, the number of ENIs per NetworkCard is 2, meaning that the only valid values for DeviceIndex is 0 or 1. Below is the example of EC2 P5 instance launch command using awscli showing NetworkCardIndex from number 0-31 and DeviceIndex as 0 for first interface and DeviceIndex as 1 for rest 31 interrfaces.

```
aws ec2 run-instances --region $REGION \
    --instance-type $INSTANCETYPE \
    --image-id $AMI --key-name $KEYNAME \
    --iam-instance-profile "Name=dlami-builder" \
    --tag-specifications "ResourceType=instance,Tags=[{Key=Name,Value=$TAG}]" \
    --network-interfaces "NetworkCardIndex=0,DeviceIndex=0,Groups=$SG,SubnetId=
$SUBNET,InterfaceType=efa" \
    "NetworkCardIndex=1,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa" \
    "NetworkCardIndex=2,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa" \
    "NetworkCardIndex=3,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa" \
    "NetworkCardIndex=4,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa" \
    "NetworkCardIndex=4,Device
```
. . .

"NetworkCardIndex=31, DeviceIndex=1, Groups=\$SG, SubnetId=\$SUBNET, InterfaceType=efa"

Kernel

• Kernel version is pinned using command:

```
echo linux-aws hold | sudo dpkg —set-selections
echo linux-headers-aws hold | sudo dpkg —set-selections
echo linux-image-aws hold | sudo dpkg —set-selections
```

• We recommend users to avoid updating their kernel version (unless due to a security patch) to ensure compatibility with installed drivers and package versions. If users still wish to update they can run the following commands to unpin their kernel versions:

```
echo linux-aws install | sudo dpkg -set-selections
echo linux-headers-aws install | sudo dpkg -set-selections
echo linux-image-aws install | sudo dpkg -set-selections
apt-get upgrade -y
```

• For each new version of DLAMI, latest available compatible kernel is used.

Release Date: 2025-02-17

AMI name: Deep Learning OSS Nvidia Driver AMI GPU PyTorch 2.5.1 (Ubuntu 22.04) 20250216

Updated

- Updated NVIDIA Container Toolkit from version 1.17.3 to version 1.17.4
 - Please see the release notes page here for more information: <u>https://github.com/NVIDIA/</u> nvidia-container-toolkit/releases/tag/v1.17.4
 - In Container Toolkit version 1.17.4, the mounting of CUDA compat libraries is now disabled. In order to ensure compatibility with multiple CUDA versions on container workflows, please ensure you update your LD_LIBRARY_PATH to include your CUDA compatibility libraries as shown in the <u>If you use a CUDA compatibility layer</u> tutorial.

Removed

 Removed user space libraries cuobj and nvdisasm provided by <u>NVIDIA CUDA toolkit</u> to address CVEs present in the NVIDIA CUDA Toolkit Security Bulletin for February 18, 2025

Release Date: 2025-01-21

AMI name: Deep Learning OSS Nvidia Driver AMI GPU PyTorch 2.5.1 (Ubuntu 22.04) 20250119

Updated

• Upgraded Nvidia driver from version 550.127.05 to 550.144.03 to address CVEs present in the NVIDIA GPU Display Driver Security Bulletin for January 2025.

Release Date: 2024-11-21

AMI name: Deep Learning OSS Nvidia Driver AMI GPU PyTorch 2.5.1 (Ubuntu 22.04) 20241121

Added

 Initial release of Deep Learning AMI GPU PyTorch 2.4.1 (Ubuntu 22.04) series. Including a conda environment pytorch complimented with NVIDIA Driver R550, CUDA=12.4.1, cuDNN=8.9.7, PyTorch NCCL=2.21.5, and EFA=1.37.0.

Fixed

- Due to a change in the Ubuntu kernel to address a defects in the Kernel Address Space Layout Randomization (KASLR) functionality, G4Dn/G5 instances are unable to properly initialize CUDA on the OSS Nvidia driver. In order to mitigate this issue, this DLAMI includes functionality that dynamically loads the proprietary driver for G4Dn and G5 instances. Please allow a brief initialization period for this loading in order to ensure that your instances are able to work properly.
 - To check the status and health of this service, you can use the following commands:

```
sudo systemctl is-active dynamic_driver_load.service
active
```

Amazon Deep Learning AMI GPU PyTorch 2.4 (Ubuntu 22.04)

For help getting started, see Getting started with DLAMI.

AMI name format

 Deep Learning OSS Nvidia Driver AMI GPU PyTorch 2.4.\${PATCH_VERSION} (Ubuntu 22.04) \${YYYY-MM-DD}

Supported EC2 instances

- Please refer to Important changes to DLAMI.
- Deep Learning with OSS Nvidia Driver supports G4dn, G5, G6, Gr6, P4, P4de, P5, P5e, P5en.

- Supported Amazon Service: EC2
- Operating System: Ubuntu 22.04
- **Compute Architecture**: x86
- Python: /opt/conda/envs/pytorch/bin/python
- NVIDIA Driver:
 - OSS Nvidia driver: 550.144.03
- NVIDIA CUDA12.1 stack:
 - CUDA, NCCL and cuDDN installation path: /usr/local/cuda-12.4/
 - Default CUDA: 12.4
 - PATH /usr/local/cuda points to /usr/local/cuda-12.4/
 - Updated below env vars:
 - LD_LIBRARY_PATH to have /usr/local/cuda/lib:/usr/local/cuda/lib64:/usr/local/cuda:/usr/ local/cuda/targets/x86_64-linux/lib
 - PATH to have /usr/local/cuda/bin/:/usr/local/cuda/include/
 - Compiled system NCCL Version present at /usr/local/cuda/: 2.21.5
 - PyTorch Compiled NCCL Version from PyTorch conda environment: 2.20.5
- NCCL Tests Location:
 - all_reduce, all_gather and reduce_scatter: /usr/local/cuda-xx.x/efa/test-cuda-xx.x/

- To run NCCL tests, LD_LIBRARY_PATH is already with updated with needed paths.
 - Common PATHs are already added to LD_LIBRARY_PATH:
 - /opt/amazon/efa/lib:/opt/amazon/openmpi/lib:/opt/aws-ofi-nccl/lib:/usr/local/lib:/usr/ lib
- LD_LIBRARY_PATH is updated with CUDA version paths
 - /usr/local/cuda/lib:/usr/local/cuda/lib64:/usr/local/cuda:/usr/local/cud/targets/x86_64linux/lib
- EFA Installer: 1.34.0
- Nvidia GDRCopy: 2.4.1
- Nvidia Transformer Engine: v1.11.0
- Amazon OFI NCCL: 1.11.0-aws
 - Installation path: /opt/aws-ofi-nccl/ . Path /opt/aws-ofi-nccl/lib is added to LD_LIBRARY_PATH.
 - Tests path for ring, message_transfer: /opt/aws-ofi-nccl/tests
 - Note: PyTorch package comes with dynamically linked Amazon OFI NCCL plugin as a conda package aws-ofi-nccl-dlc package as well and PyTorch will use that package instead of system Amazon OFI NCCL.
- Amazon CLI v2 as aws2 and Amazon CLI v1 as aws
- EBS volume type: gp3
- Python version: 3.11
- Query AMI-ID with SSM Parameter (example Region is us-east-1):
 - OSS Nvidia Driver:

- Query AMI-ID with AWSCLI (example Region is us-east-1):
 - OSS Nvidia Driver:

```
aws ec2 describe-images --region us-east-1 \
    --owners amazon \
```

```
--filters 'Name=name,Values=Deep Learning OSS Nvidia Driver AMI GPU PyTorch
2.4.? (Ubuntu 22.04) ???????' 'Name=state,Values=available' \
    --query 'reverse(sort_by(Images, &CreationDate))[:1].ImageId' \
    --output text
```

Notices

P5/P5e instances

 DeviceIndex is unique to each NetworkCard, and must be a non-negative integer less than the limit of ENIs per NetworkCard. On P5, the number of ENIs per NetworkCard is 2, meaning that the only valid values for DeviceIndex is 0 or 1. Below is the example of EC2 P5 instance launch command using awscli showing NetworkCardIndex from number 0-31 and DeviceIndex as 0 for first interface and DeviceIndex as 1 for rest 31 interrfaces.

```
aws ec2 run-instances --region $REGION \
    --instance-type $INSTANCETYPE \
    --image-id $AMI --key-name $KEYNAME \
    --iam-instance-profile "Name=dlami-builder" \
    --tag-specifications "ResourceType=instance,Tags=[{Key=Name,Value=$TAG}]" \
    --network-interfaces "NetworkCardIndex=0, DeviceIndex=0, Groups=$SG, SubnetId=
$SUBNET, InterfaceType=efa" ∖
      "NetworkCardIndex=1, DeviceIndex=1, Groups=$SG, SubnetId=$SUBNET, InterfaceType=efa"
 /
      "NetworkCardIndex=2,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa"
 /
      "NetworkCardIndex=3,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa"
 \
      "NetworkCardIndex=4,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa"
 \
      . . .
      "NetworkCardIndex=31,DeviceIndex=1,Groups=$SG,SubnetId=$SUBNET,InterfaceType=efa"
```

Release Date: 2025-02-17

AMI name: Deep Learning OSS Nvidia Driver AMI GPU PyTorch 2.4.1 (Ubuntu 22.04) 20250216

Updated

Updated NVIDIA Container Toolkit from version 1.17.3 to version 1.17.4

- Please see the release notes page here for more information: <u>https://github.com/NVIDIA/</u> <u>nvidia-container-toolkit/releases/tag/v1.17.4</u>
- In Container Toolkit version 1.17.4, the mounting of CUDA compat libraries is now disabled. In order to ensure compatibility with multiple CUDA versions on container workflows, please ensure you update your LD_LIBRARY_PATH to include your CUDA compatibility libraries as shown in the <u>If you use a CUDA compatibility layer</u> tutorial.

Release Date: 2025-01-21

AMI name: Deep Learning OSS Nvidia Driver AMI GPU PyTorch 2.4.1 (Ubuntu 22.04) 20250119

Updated

 Upgraded Nvidia driver from version 550.127.05 to 550.144.03 to address CVE's present in the NVIDIA GPU Display Driver Security Bulletin for January 2025.

Release Date: 2024-11-18

AMI name: Deep Learning OSS Nvidia Driver AMI GPU PyTorch 2.4.1 (Ubuntu 22.04) 20241116

Fixed

- Due to a change in the Ubuntu kernel to address a defects in the Kernel Address Space Layout Randomization (KASLR) functionality, G4Dn/G5 instances are unable to properly initialize CUDA on the OSS Nvidia driver. In order to mitigate this issue, this DLAMI includes functionality that dynamically loads the proprietary driver for G4Dn and G5 instances. Please allow a brief initialization period for this loading in order to ensure that your instances are able to work properly.
 - To check the status and health of this service, you can use the following commands:

```
sudo systemctl is-active dynamic_driver_load.service
active
```

Release Date: 2024-10-16

AMI name: Deep Learning OSS Nvidia Driver AMI GPU PyTorch 2.4.1 (Ubuntu 22.04) 20241016

Added

 Added Nvidia TransformerEngine v1.11.0 for accelerating Transformer models (For more details, please refer to <u>https://docs.nvidia.com/deeplearning/transformer-engine/user-guide/</u> index.html)

Release Date: 2024-09-30

AMI name: Deep Learning OSS Nvidia Driver AMI GPU PyTorch 2.4.1 (Ubuntu 22.04) 20240929

Updated

 Upgraded Nvidia Container Toolkit from version 1.16.1 to 1.16.2, addressing the security vulnerability <u>CVE-2024-0133</u>.

Release Date: 2024-09-26

AMI name: Deep Learning OSS Nvidia Driver AMI GPU PyTorch 2.4.1 (Ubuntu 22.04) 20240925

Added

 Initial release of Deep Learning AMI GPU PyTorch 2.4.1 (Ubuntu 22.04) series. Including a conda environment pytorch complimented with NVIDIA Driver R550, CUDA=12.4.1, cuDNN=8.9.7, PyTorch NCCL=2.20.5, and EFA=1.34.0.

ARM64 PyTorch DLAMI Release Notes

Below Are the Release notes for ARM64 PyTorch DLAMIs:

GPU

- Amazon Deep Learning ARM64 AMI GPU PyTorch 2.7 (Amazon Linux 2023)
- Amazon Deep Learning ARM64 AMI GPU PyTorch 2.7 (Ubuntu 22.04)
- Amazon Deep Learning ARM64 AMI GPU PyTorch 2.6 (Amazon Linux 2023)
- Amazon Deep Learning ARM64 AMI GPU PyTorch 2.6 (Ubuntu 22.04)
- Amazon Deep Learning ARM64 AMI GPU PyTorch 2.5 (Ubuntu 22.04)
- Amazon Deep Learning ARM64 AMI GPU PyTorch 2.4 (Ubuntu 22.04)

Amazon Neuron

• Refer to the Neuron DLAMI User Guide

Amazon Deep Learning ARM64 AMI OSS GPU PyTorch 2.7 (Amazon Linux 2023)

For help getting started, see <u>Getting started with DLAMI</u>.

AMI name format

 Deep Learning ARM64 AMI OSS Nvidia Driver GPU PyTorch 2.7 (Amazon Linux 2023) \${YYYY-MM-DD}

Supported EC2 instances

• G5g

- Supported Amazon Service: Amazon EC2
- Operating System: Amazon Linux 2023
- Compute Architecture: ARM64
- Linux Kernel: 6.12
- NVIDIA Driver: 570.133.20
- NVIDIA CUDA 12.8 stack:
 - CUDA, NCCL and cuDDN installation directories: /usr/local/cuda-12.8/
 - Default CUDA: 12.8
 - PATH /usr/local/cuda points to CUDA 12.8
 - Updated below env vars:
 - LD_LIBRARY_PATH to have /usr/local/cuda/lib:/usr/local/cuda/lib64:/usr/local/cuda:/usr/ local/cuda/targets/sbsa-
 - PATH to have /usr/local/cuda/bin/:/usr/local/cuda/include/
- Amazon CLI v2 at /usr/local/bin/aws
- EBS volume type: gp3
- Nvidia container toolkit: 1.17.7

- Version command: nvidia-container-cli -V
- Docker: 25.0.8
- **Python**: /usr/bin/python3.12
- Query AMI-ID with SSM Parameter (example region is us-east-1):

```
aws ssm get-parameter --region us-east-1 \
    --name /aws/service/deeplearning/ami/arm64/oss-nvidia-driver-gpu-pytorch-2.7-
amazon-linux-2023/latest/ami-id \
    --query "Parameter.Value" \
    --output text
```

• Query AMI-ID with AWSCLI (example region is us-east-1):

```
aws ec2 describe-images --region us-east-1 --owners amazon --filters
'Name=name,Values=Deep Learning ARM64 AMI OSS Nvidia Driver GPU PyTorch 2.7 (Amazon
Linux 2023) ???????' 'Name=state,Values=available' --query 'reverse(sort_by(Images,
&CreationDate))[:1].ImageId' --output text
```

Notices

PyTorch Deprecation of Anaconda Channel

Starting with PyTorch 2.6, PyTorch has deprecated support for Conda (see <u>official announcement</u>). As a result, PyTorch 2.6 and above will move to using Python Virtual Environments. To activate the PyTorch venv, please use source /opt/pytorch/bin/activate

Release Date: 2025-05-22

AMI name: Deep Learning ARM64 AMI OSS Nvidia Driver GPU PyTorch 2.7 (Amazon Linux 2023) 20250521

Added

Initial release of Deep Learning ARM64 AMI OSS Nvidia Driver GPU PyTorch 2.7 (Amazon Linux 2023) series. Including a Python virtual environment pytorch (source /opt/pytorch/bin/activate) complimented with NVIDIA Driver R570, CUDA=12.8, cuDNN=9.10, and PyTorch NCCL=2.26.2

Amazon Deep Learning ARM64 AMI OSS GPU PyTorch 2.7 (Ubuntu 22.04)

For help getting started, see Getting started with DLAMI.

AMI name format

• Deep Learning ARM64 AMI OSS Nvidia Driver GPU PyTorch 2.7 (Ubuntu 22.04) \${YYYY-MM-DD}

Supported EC2 instances

• G5g

- Supported Amazon Service: Amazon EC2
- Operating System: Ubuntu 22.04
- Compute Architecture: ARM64
- Linux Kernel: 6.8
- NVIDIA Driver: 570.133.20
- NVIDIA CUDA 12.8 stack:
 - CUDA, NCCL and cuDDN installation directories: /usr/local/cuda-12.8/
 - Default CUDA: 12.8
 - PATH /usr/local/cuda points to CUDA 12.8
 - Updated below env vars:
 - LD_LIBRARY_PATH to have /usr/local/cuda/lib:/usr/local/cuda/lib64:/usr/local/cuda:/usr/ local/cuda/targets/sbsa-
 - PATH to have /usr/local/cuda/bin/:/usr/local/cuda/include/
- Amazon CLI v2 at /usr/local/bin/aws
- EBS volume type: gp3
- Nvidia container toolkit: 1.17.7
 - Version command: nvidia-container-cli -V
- Docker: 28.2.2
- **Python**: /usr/bin/python3.12
- Query AMI-ID with SSM Parameter (example region is us-east-1):

```
aws ssm get-parameter --region us-east-1 \
    --name /aws/service/deeplearning/ami/arm64/oss-nvidia-driver-gpu-pytorch-2.7-
ubuntu-22.04/latest/ami-id \
    --query "Parameter.Value" \
    --output text
```

• Query AMI-ID with AWSCLI (example region is us-east-1):

```
aws ec2 describe-images --region us-east-1 --owners amazon --filters
'Name=name,Values=Deep Learning ARM64 AMI OSS Nvidia Driver GPU PyTorch 2.7 (Ubuntu
22.04) ???????' 'Name=state,Values=available' --query 'reverse(sort_by(Images,
&CreationDate))[:1].ImageId' --output text
```

Notices

PyTorch Deprecation of Anaconda Channel

Starting with PyTorch 2.6, PyTorch has deprecated support for Conda (see <u>official announcement</u>). As a result, PyTorch 2.6 and above will move to using Python Virtual Environments. To activate the PyTorch venv, please use source /opt/pytorch/bin/activate

Release Date: 2025-05-22

AMI name: Deep Learning ARM64 AMI OSS Nvidia Driver GPU PyTorch 2.7 (Ubuntu 22.04) 20250521

Added

 Initial release of Deep Learning ARM64 AMI OSS Nvidia Driver GPU PyTorch 2.7 (Ubuntu 22.04) series. Including a Python virtual environment pytorch (source /opt/pytorch/bin/activate) complimented with NVIDIA Driver R570, CUDA=12.8, cuDNN=9.10, and PyTorch NCCL=2.26.2

Amazon Deep Learning ARM64 AMI GPU PyTorch 2.6 (Amazon Linux 2023)

For help getting started, see Getting started with DLAMI.

AMI name format

 Deep Learning ARM64 AMI OSS Nvidia Driver GPU PyTorch 2.6.\${PATCH-VERSION} (Amazon Linux 2023) \${YYYY-MM-DD}

Supported EC2 instances

• G5g

- Supported Amazon Service: EC2
- Operating System: Amazon Linux 2023
- Compute Architecture: ARM64
- Python: /opt/pytorch/bin/python
- Python version: 3.12
- NVIDIA Driver:
 - OSS Nvidia driver:570.86.15
- NVIDIA CUDA12.6 stack:
 - CUDA, NCCL and cuDDN installation path: /usr/local/cuda-12.6/
 - Default CUDA: 12.6
 - PATH /usr/local/cuda points to /usr/local/cuda-12.6/
 - Updated below env vars:
 - LD_LIBRARY_PATH to have /usr/local/cuda/lib:/usr/local/cuda/lib64:/usr/local/cuda:/ usr/local/cuda/targets/sbsa-linux/lib:/usr/local/cuda/nvvm/lib64:/usr/local/cuda/extras/ CUPTI/lib64
 - PATH to have /usr/local/cuda/bin/:/usr/local/cuda/include/
 - Compiled system NCCL Version present at /usr/local/cuda/: 2.24.3
 - PyTorch Compiled NCCL Version from PyTorch conda environment: 2.21.5
- Amazon CLI v2 as aws2 and Amazon CLI v1 as aws
- EBS volume type: gp3
- Query AMI-ID with SSM Parameter (example Region is us-east-1):

```
aws ssm get-parameter --region us-east-1 \
    --name /aws/service/deeplearning/ami/arm64/oss-nvidia-driver-gpu-pytorch-2.6-
amazon-linux-2023/latest/ami-id \
    --query "Parameter.Value" \
    --output text
```

• Query AMI-ID with AWSCLI (example Region is us-east-1):

```
aws ec2 describe-images --region us-east-1 \
    --owners amazon \
    --filters 'Name=name,Values=Deep Learning ARM64 AMI OSS Nvidia Driver GPU PyTorch
2.6.? (Amazon Linux 2023) ???????' 'Name=state,Values=available' \
    --query 'reverse(sort_by(Images, &CreationDate))[:1].ImageId' \
    --output text
```

Notices

PyTorch Deprecation of Anaconda Channel

Starting with PyTorch 2.6, PyTorch has deprecated support for Conda (see <u>official announcement</u>). As a result, PyTorch 2.6 and above will move to using Python Virtual Environments. To activate the PyTorch venv, please use source /opt/pytorch/bin/activate

Release Date: 2025-02-21

AMI name: Deep Learning ARM64 AMI OSS Nvidia Driver GPU PyTorch 2.6.0 (Amazon Linux 2023) 20250221

Added

 Initial release of Deep Learning ARM64 AMI OSS Nvidia Driver GPU PyTorch 2.6.0 (Amazon Linux 2023) series. Including a Python virtual environment pytorch (source /opt/pytorch/ bin/activate/), complimented with NVIDIA Driver R570, CUDA=12.6, cuDNN=9.7, PyTorch NCCL=2.21.5.

Amazon Deep Learning ARM64 AMI GPU PyTorch 2.6 (Ubuntu 22.04)

For help getting started, see Getting started with DLAMI.

AMI name format

Deep Learning ARM64 AMI OSS NVIDIA Driver GPU PyTorch 2.6.\${PATCH-VERSION} (Ubuntu 22.04) \${YYYY-MM-DD}

Supported EC2 instances

• G5g

- Supported Amazon Service: Amazon EC2
- Operating System: Ubuntu 22.04
- Compute Architecture: ARM64
- Python: /opt/pytorch/bin/python
- Python version: 3.12
- NVIDIA Driver:
 - OSS Nvidia driver: 570.86.15
- NVIDIA CUDA12.6 stack:
 - CUDA, NCCL and cuDDN installation path: /usr/local/cuda-12.6/
 - Default CUDA: 12.6
 - PATH /usr/local/cuda points to /usr/local/cuda-12.6/
 - Updated below env vars:
 - LD_LIBRARY_PATH to have /usr/local/cuda/lib:/usr/local/cuda/lib64:/usr/local/cuda:/ usr/local/cuda/targets/sbsa-linux/lib:/usr/local/cuda/nvvm/lib64:/usr/local/cuda/extras/ CUPTI/lib64
 - PATH to have /usr/local/cuda/bin/:/usr/local/cuda/include/
 - Compiled system NCCL Version present at /usr/local/cuda/: 2.24.3
 - PyTorch Compiled NCCL Version from PyTorch venv environment: 2.21.5
- Amazon CLI v2 as aws2 and Amazon CLI v1 as aws
- EBS volume type: gp3
- Query AMI-ID with SSM Parameter (example region is us-east-1):

```
aws ssm get-parameter --region us-east-1 \
    --name /aws/service/deeplearning/ami/arm64/oss-nvidia-driver-gpu-pytorch-2.6-
ubuntu-22.04/latest/ami-id \
    --query "Parameter.Value" \
    --output text
```

• Query AMI-ID with AWSCLI (example region is us-east-1):

```
aws ec2 describe-images --region us-east-1 \
    --owners amazon --filters 'Name=name,Values=Deep Learning ARM64 AMI OSS Nvidia
Driver GPU PyTorch 2.6.? (Ubuntu 22.04) ???????' 'Name=state,Values=available' \
    --query 'reverse(sort_by(Images, &CreationDate))[:1].ImageId' \
    --output text
```

Notices

PyTorch Deprecation of Anaconda Channel

Starting with PyTorch 2.6, PyTorch has deprecated support for Conda (see <u>official announcement</u>). As a result, PyTorch 2.6 and above will move to using Python Virtual Environments. To activate the PyTorch venv, please use source /opt/pytorch/bin/activate

Release Date: 2025-02-21

AMI name: Deep Learning ARM64 AMI OSS NVIDIA Driver GPU PyTorch 2.6.0 (Ubuntu 22.04) 20250221

Added

Initial release of Deep Learning ARM64 AMI OSS Nvidia Driver GPU PyTorch 2.6.0 (Ubuntu 22.04) series. Including a Python virtual environment pytorch (source /opt/pytorch/bin/activate), complimented with NVIDIA Driver R570, CUDA=12.6, cuDNN=9.7, PyTorch NCCL=2.21.5.

Amazon Deep Learning ARM64 AMI GPU PyTorch 2.5 (Ubuntu 22.04)

For help getting started, see Getting started with DLAMI.

AMI name format

Deep Learning ARM64 AMI OSS Nvidia Driver GPU PyTorch 2.5.\${PATCH-VERSION} (Ubuntu 22.04) \${YYYY-MM-DD}

Supported EC2 instances

• G5g

The AMI includes the following:

- Supported Amazon Service: Amazon EC2
- Operating System: Ubuntu 22.04
- Compute Architecture: ARM64
- Python: /opt/conda/envs/pytorch/bin/python
- Python version: 3.11
- NVIDIA Driver:
 - OSS Nvidia driver: 550.144.03
- NVIDIA CUDA12.4 stack:
 - CUDA, NCCL and cuDDN installation path: /usr/local/cuda-12.4/
 - Default CUDA: 12.4
 - PATH /usr/local/cuda points to /usr/local/cuda-12.4/
 - Updated below env vars:
 - LD_LIBRARY_PATH to have /usr/local/cuda/lib:/usr/local/cuda/lib64:/usr/local/cuda:/ usr/local/cuda/targets/sbsa-linux/lib:/usr/local/cuda/nvvm/lib64:/usr/local/cuda/extras/ CUPTI/lib64
 - PATH to have /usr/local/cuda/bin/:/usr/local/cuda/include/
 - Compiled system NCCL Version present at /usr/local/cuda/: 2.21.5
 - PyTorch Compiled NCCL Version from PyTorch conda environment: 2.21.5
- Amazon CLI v2 as aws2 and Amazon CLI v1 as aws
- EBS volume type: gp3
- Query AMI-ID with SSM Parameter (example Region is us-east-1):

```
aws ssm get-parameter --region us-east-1 \setminus
    --name /aws/service/deeplearning/ami/arm64/oss-nvidia-driver-gpu-pytorch-2.5-
ubuntu-22.04/latest/ami-id \
    --query "Parameter.Value" \
    --output text
```

Query AMI-ID with AWSCLI (example Region is us-east-1):

```
aws ec2 describe-images --region us-east-1 \
    --owners amazon --filters \
```

PyTorch DLAMIs

```
'Name=name,Values=Deep Learning ARM64 AMI OSS Nvidia Driver GPU PyTorch 2.5.?
(Ubuntu 22.04) ??????' 'Name=state,Values=available' \
    --query 'reverse(sort_by(Images, &CreationDate))[:1].ImageId' \
    --output text
```

Release Date: 2025-02-17

AMI name: Deep Learning ARM64 AMI OSS Nvidia Driver GPU PyTorch 2.5.1 (Ubuntu 22.04) 20250215

Updated

- Updated NVIDIA Container Toolkit from version 1.17.3 to version 1.17.4
 - Please see the release notes page here for more information: <u>https://github.com/NVIDIA/</u> nvidia-container-toolkit/releases/tag/v1.17.4
 - In Container Toolkit version 1.17.4, the mounting of CUDA compat libraries is now disabled. In order to ensure compatibility with multiple CUDA versions on container workflows, please ensure you update your LD_LIBRARY_PATH to include your CUDA compatibility libraries as shown in the <u>If you use a CUDA compatibility layer</u> tutorial.

Removed

 Removed user space libraries cuobj and nvdisasm provided by <u>NVIDIA CUDA toolkit</u> to address CVEs present in the <u>NVIDIA CUDA Toolkit Security Bulletin for February 18, 2025</u>

Release Date: 2025-01-21

AMI name: Deep Learning ARM64 AMI OSS Nvidia Driver GPU PyTorch 2.5.1 (Ubuntu 22.04) 20250117

Updated

 Upgraded Nvidia driver from version 550.127.05 to 550.144.03 to address CVEs present in the NVIDIA GPU Display Driver Security Bulletin for January 2025.

Release Date: 2024-11-22

AMI name: Deep Learning ARM64 AMI OSS Nvidia Driver GPU PyTorch 2.5.1 (Ubuntu 22.04) 20241122

Added

Initial release of Deep Learning ARM64 AMI OSS Nvidia Driver GPU PyTorch 2.5.1 (Ubuntu 22.04) series. Including a conda environment pytorch complimented with NVIDIA Driver R550, CUDA=12.4, cuDNN=8.9.7, PyTorch NCCL=2.21.5.

Amazon Deep Learning ARM64 AMI GPU PyTorch 2.4 (Ubuntu 22.04)

For help getting started, see Getting started with DLAMI.

AMI name format

Deep Learning ARM64 AMI OSS Nvidia Driver GPU PyTorch 2.4.\${PATCH-VERSION} (Ubuntu 22.04) \${YYYY-MM-DD}

Supported EC2 instances

• G5g

- Supported Amazon Service: Amazon EC2
- Operating System: Ubuntu 22.04
- Compute Architecture: ARM64
- Python: /opt/conda/envs/pytorch/bin/python
- Python version: 3.11
- NVIDIA Driver:
 - OSS Nvidia driver: 550.144.03
- NVIDIA CUDA12.1 stack:
 - CUDA, NCCL and cuDDN installation path: /usr/local/cuda-12.4/
 - Default CUDA: 12.4

- PATH /usr/local/cuda points to /usr/local/cuda-12.4/
- Updated below env vars:
 - LD_LIBRARY_PATH to have /usr/local/cuda/lib:/usr/local/cuda/lib64:/usr/local/cuda:/ usr/local/cuda/targets/sbsa-linux/lib:/usr/local/cuda/nvvm/lib64:/usr/local/cuda/extras/ CUPTI/lib64
 - PATH to have /usr/local/cuda/bin/:/usr/local/cuda/include/
- Compiled system NCCL Version present at /usr/local/cuda/: 2.21.5
- PyTorch Compiled NCCL Version from PyTorch conda environment: 2.20.5
- Amazon CLI v2 as aws2 and Amazon CLI v1 as aws
- EBS volume type: gp3
- Query AMI-ID with SSM Parameter (example Region is us-east-1):

```
aws ssm get-parameter --region us-east-1 \
    --name /aws/service/deeplearning/ami/arm64/oss-nvidia-driver-gpu-pytorch-2.4-
ubuntu-22.04/latest/ami-id \
    --query "Parameter.Value" \
    --output text
```

• Query AMI-ID with AWSCLI (example Region is us-east-1):

```
aws ec2 describe-images --region us-east-1 \
    --owners amazon \
    --filters 'Name=name,Values=Deep Learning ARM64 AMI OSS Nvidia Driver GPU PyTorch
2.4.? (Ubuntu 22.04) ???????' 'Name=state,Values=available' \
    --query 'reverse(sort_by(Images, &CreationDate))[:1].ImageId' \
    --output text
```

Release Date: 2025-02-17

AMI name: Deep Learning ARM64 AMI OSS Nvidia Driver GPU PyTorch 2.4.0 (Ubuntu 22.04) 20250215

Updated

- Updated NVIDIA Container Toolkit from version 1.17.3 to version 1.17.4
 - Please see the release notes page here for more information: <u>https://github.com/NVIDIA/</u> nvidia-container-toolkit/releases/tag/v1.17.4

 In Container Toolkit version 1.17.4, the mounting of CUDA compat libraries is now disabled. In order to ensure compatibility with multiple CUDA versions on container workflows, please ensure you update your LD_LIBRARY_PATH to include your CUDA compatibility libraries as shown in the If you use a CUDA compatibility layer tutorial.

Removed

 Removed user space libraries cuobj and nvdisasm provided by <u>NVIDIA CUDA toolkit</u> to address CVEs present in the <u>NVIDIA CUDA Toolkit Security Bulletin for February 18, 2025</u>

Release Date: 2025-01-21

AMI name: Deep Learning ARM64 AMI OSS Nvidia Driver GPU PyTorch 2.4.0 (Ubuntu 22.04) 20250117

Updated

• Upgraded Nvidia driver from version 550.127.05 to 550.144.03 to address CVEs present in the NVIDIA GPU Display Driver Security Bulletin for January 2025.

Release Date: 2024-09-30

AMI name: Deep Learning ARM64 AMI OSS Nvidia Driver GPU PyTorch 2.4.0 (Ubuntu 22.04) 20240927

Updated

• Upgraded Nvidia Container Toolkit from version 1.16.1 to 1.16.2, addressing the security vulnerability CVE-2024-0133.

Release Date: 2024-09-26

AMI name: Deep Learning ARM64 AMI OSS Nvidia Driver GPU PyTorch 2.4.0 (Ubuntu 22.04) 20240926

Added

Initial release of Deep Learning ARM64 AMI OSS Nvidia Driver GPU PyTorch 2.4.0 (Ubuntu 22.04) series. Including a conda environment pytorch complimented with NVIDIA Driver R550, CUDA=12.4, cuDNN=8.9.7, PyTorch NCCL=2.20.5.

TensorFlow DLAMIs

TensorFlow DLAMI Release Notes

X86 TensorFlow DLAMI Release Notes

X86 TensorFlow DLAMI Release Notes

Below are the release notes for X86 TensorFlow DLAMI:

GPU

- Amazon Deep Learning AMI GPU TensorFlow 2.18 (Amazon Linux 2023)
- Amazon Deep Learning AMI GPU TensorFlow 2.18 (Ubuntu 22.04)
- Amazon Deep Learning AMI GPU TensorFlow 2.17 (Ubuntu 22.04)

Amazon Neuron

• Refer to Neuron DLAMI Guide.

Deep Learning AMI GPU TensorFlow 2.18 (Amazon Linux 2023)

For help getting started, see Getting started with DLAMI.

AMI name format

 Deep Learning OSS Nvidia Driver AMI GPU TensorFlow 2.18 (Amazon Linux 2023) \${YYYY-MM-DD}

Supported EC2 instances

• Deep Learning with OSS Nvidia Driver supports G4dn, G5, G6, Gr6, G6e, P4d, P4de, P5, P5e, P5en

- Supported Amazon Service: Amazon EC2
- Operating System: Amazon Linux 2023
- Compute Architecture: x86
- Python: /opt/tensorflow/bin/python3.12
- TensorFlow version: 2.18
- NVIDIA Driver:
 - OSS Nvidia driver: 560.35.03
- NVIDIA CUDA12 stack:
 - CUDA, NCCL and cuDDN installation path: /usr/local/cuda-12.5/
- EFA Installer: 1.37.0
- Amazon CLI v2 as aws2 and Amazon CLI v1 as aws
- EBS volume type: gp3
- Query AMI-ID with SSM Parameter (example Region is us-east-1):
 - OSS Nvidia Driver:

```
aws ssm get-parameter --region us-east-1 \
    --name /aws/service/deeplearning/ami/x86_64/oss-nvidia-driver-gpu-
tensorflow-2.18-amazon-linux-2023/latest/ami-id \
    --query "Parameter.Value" \
    --output text
```

- Query AMI-ID with AWSCLI (example Region is us-east-1):
 - OSS Nvidia Driver:

```
aws ec2 describe-images --region us-east-1 \
    --owners amazon \
    --filters 'Name=name,Values=Deep Learning OSS Nvidia Driver AMI GPU TensorFlow
2.18 (Amazon Linux 2023) ???????' 'Name=state,Values=available' \
    --query 'reverse(sort_by(Images, &CreationDate))[:1].ImageId' \
    --output text
```

Release Date: 2025-02-17

AMI name: Deep Learning OSS Nvidia Driver AMI GPU TensorFlow 2.18 (Amazon Linux 2023) 20241215

Updated

- Updated NVIDIA Container Toolkit from version 1.17.3 to version 1.17.4
 - Please see the release notes page here for more information: <u>https://github.com/NVIDIA/</u> nvidia-container-toolkit/releases/tag/v1.17.4
 - In Container Toolkit version 1.17.4, the mounting of CUDA compat libraries is now disabled. In order to ensure compatibility with multiple CUDA versions on container workflows, please ensure you update your LD_LIBRARY_PATH to include your CUDA compatibility libraries as shown in the <u>If you use a CUDA compatibility layer</u> tutorial.

Removed

 Removed user space libraries cuobj and nvdisasm provided by <u>NVIDIA CUDA toolkit</u> to address CVEs present in the <u>NVIDIA CUDA Toolkit Security Bulletin for February 18, 2025</u>

Release Date: 2024-12-09

AMI name: Deep Learning OSS Nvidia Driver AMI GPU TensorFlow 2.18 (Amazon Linux 2023) 20241206

Added

- Initial release of the Deep Learning OSS Nvidia Driver AMI GPU TensorFlow 2.18 (Amazon Linux 2023) series.
 - Software Includes the Following:
 - "nvidia-driver=560.35.03"
 - "fabric-manager=560.35.03"
 - "cuda=12.5"
 - "cudnn=9.5.1"
 - "efa=1.37.0"
 - "nccl=2.23.4"
 - "aws-nccl-ofi-plugin=v1.13.0-aws"

- Tensorflow virtual environment (activation command source /opt/tensorflow/bin/activate) includes the following:
 - "tensorflow=2.18.0"

Deep Learning AMI GPU TensorFlow 2.18 (Ubuntu 22.04)

For help getting started, see <u>Getting started with DLAMI</u>.

AMI name format

• Deep Learning OSS Nvidia Driver AMI GPU TensorFlow 2.18 (Ubuntu 22.04) \${YYYY-MM-DD}

Supported EC2 instances

• Deep Learning with OSS Nvidia Driver supports G4dn, G5, G6, Gr6, G6e, P4d, P4de, P5, P5e, P5en.

The AMI includes the following:

- Supported Amazon Service: Amazon EC2
- Operating System: Ubuntu 22.04
- Compute Architecture: x86
- Python: /opt/tensorflow/bin/python3.12
- TensorFlow version: 2.18
- NVIDIA Driver:
 - OSS Nvidia driver: 550.144.03
- NVIDIA CUDA12 stack:
 - CUDA, NCCL and cuDDN installation path: /usr/local/cuda-12.5/
- **EFA Installer**: 1.37.0
- Amazon CLI v2 as aws2 and Amazon CLI v1 as aws
- EBS volume type: gp3
- Query AMI-ID with SSM Parameter (example Region is us-east-1):
 - OSS Nvidia Driver:

aws ssm get-parameter --region $us-east-1 \setminus$

```
--name /aws/service/deeplearning/ami/x86_64/oss-nvidia-driver-gpu-
tensorflow-2.18-ubuntu-22.04/latest/ami-id \
    --query "Parameter.Value" \
    --output text
```

- Query AMI-ID with AWSCLI (example Region is us-east-1):
 - OSS Nvidia Driver:

```
aws ec2 describe-images --region us-east-1 \
    --owners amazon --filters 'Name=name,Values=Deep Learning OSS Nvidia Driver AMI
GPU TensorFlow 2.18 (Ubuntu 22.04) ??????' 'Name=state,Values=available' \
    --query 'reverse(sort_by(Images, &CreationDate))[:1].ImageId' \
    --output text
```

Release Date: 2025-02-17

AMI name: Deep Learning OSS Nvidia Driver AMI GPU TensorFlow 2.18 (Ubuntu 22.04) 20250215

Updated

- Updated NVIDIA Container Toolkit from version 1.17.3 to version 1.17.4
 - Please see the release notes page here for more information: <u>https://github.com/NVIDIA/</u> nvidia-container-toolkit/releases/tag/v1.17.4
 - In Container Toolkit version 1.17.4, the mounting of CUDA compat libraries is now disabled. In order to ensure compatibility with multiple CUDA versions on container workflows, please ensure you update your LD_LIBRARY_PATH to include your CUDA compatibility libraries as shown in the <u>If you use a CUDA compatibility layer</u> tutorial.

Removed

 Removed user space libraries cuobj and nvdisasm provided by <u>NVIDIA CUDA toolkit</u> to address CVEs present in the <u>NVIDIA CUDA Toolkit Security Bulletin for February 18, 2025</u>

Release Date: 2025-01-20

AMI name: Deep Learning OSS Nvidia Driver AMI GPU TensorFlow 2.18 (Ubuntu 22.04) 20250118

Updated

• Upgraded Nvidia driver from version 550.90.07 to 550.127.05 to address CVEs present in the NVIDIA GPU Display Driver Security Bulletin for January 2025

Release Date: 2024-12-09

AMI name: Deep Learning OSS Nvidia Driver AMI GPU TensorFlow 2.18 (Ubuntu 22.04) 20241206

Added

- Initial release of the Deep Learning OSS Nvidia Driver AMI GPU TensorFlow 2.18 (Ubuntu 22.04) series.
 - Software Includes the Following:
 - "nvidia-driver=550.127.05"
 - "fabric-manager=550.127.05"
 - "cuda=12.5"
 - "cudnn=9.5.1"
 - "efa=1.37.0"
 - "nccl=2.23.4"
 - "aws-nccl-ofi-plugin=v1.13.0-aws"
- Tensorflow virtual environment (activation command source /opt/tensorflow/bin/activate) includes the following:
 - "tensorflow=2.18.0"

Fixed

- Due to a change in the Ubuntu kernel to address a defects in the Kernel Address Space Layout Randomization (KASLR) functionality, G4Dn/G5 instances are unable to properly initialize CUDA on the OSS Nvidia driver. In order to mitigate this issue, this DLAMI includes functionality that dynamically loads the proprietary driver for G4Dn and G5 instances. Please allow a brief initialization period for this loading in order to ensure that your instances are able to work properly.
 - To check the status and health of this service, you can use the following commands:

sudo systemctl is-active dynamic_driver_load.service
active

Deep Learning AMI GPU TensorFlow 2.17 (Ubuntu 22.04)

For help getting started, see Getting started with DLAMI.

AMI name format

• Deep Learning OSS Nvidia Driver AMI GPU TensorFlow 2.17 (Ubuntu 22.04) \${YYYY-MM-DD}

Supported EC2 instances

• Deep Learning with OSS Nvidia Driver supports G4dn, G5, G6, Gr6, G6e, P4d, P4de, P5, P5e. P5en.

- Supported Amazon Service: Amazon EC2
- Operating System: Ubuntu 22.04
- Compute Architecture: x86
- Python: /opt/tensorflow/bin/python3.12
- TensorFlow version: 2.17
- NVIDIA Driver:
 - OSS Nvidia driver: 550.144.03
- NVIDIA CUDA12 stack:
 - CUDA, NCCL and cuDDN installation path: /usr/local/cuda-12.3/
- EFA Installer: 1.34.0
- Amazon CLI v2 as aws2 and Amazon CLI v1 as aws
- EBS volume type: gp3
- Query AMI-ID with SSM Parameter (example Region is us-east-1):
 - OSS Nvidia Driver:

```
--query "Parameter.Value" \
--output text
```

- Query AMI-ID with AWSCLI (example Region is us-east-1):
 - OSS Nvidia Driver:

```
aws ec2 describe-images --region us-east-1 \
    --owners amazon --filters 'Name=name,Values=Deep Learning OSS Nvidia Driver AMI
GPU TensorFlow 2.17 (Ubuntu 22.04) ??????' 'Name=state,Values=available' \
    --query 'reverse(sort_by(Images, &CreationDate))[:1].ImageId' \
    --output text
```

Release Date: 2025-02-17

AMI name: Deep Learning OSS Nvidia Driver AMI GPU TensorFlow 2.17 (Ubuntu 22.04) 20250215

Updated

- Updated NVIDIA Container Toolkit from version 1.17.3 to version 1.17.4
 - Please see the release notes page here for more information: <u>https://github.com/NVIDIA/</u> nvidia-container-toolkit/releases/tag/v1.17.4
 - In Container Toolkit version 1.17.4, the mounting of CUDA compat libraries is now disabled. In order to ensure compatibility with multiple CUDA versions on container workflows, please ensure you update your LD_LIBRARY_PATH to include your CUDA compatibility libraries as shown in the <u>If you use a CUDA compatibility layer</u> tutorial.

Removed

 Removed user space libraries cuobj and nvdisasm provided by <u>NVIDIA CUDA toolkit</u> to address CVEs present in the <u>NVIDIA CUDA Toolkit Security Bulletin for February 18, 2025</u>

Release Date: 2025-01-20

AMI name: Deep Learning OSS Nvidia Driver AMI GPU TensorFlow 2.17 (Ubuntu 22.04) 20250118

Updated

 Upgraded Nvidia driver from version 550.127.05 to 550.144.03 to address CVEs present in the NVIDIA GPU Display Driver Security Bulletin for January 2025

Version 2.17.1

Release date: 2024-11-18

AMI name: Deep Learning OSS Nvidia Driver AMI GPU TensorFlow 2.17 (Ubuntu 22.04) 20241115

Fixed

- Due to a change in the Ubuntu kernel to address a defects in the Kernel Address Space Layout Randomization (KASLR) functionality, G4Dn/G5 instances are unable to properly initialize CUDA on the OSS Nvidia driver. In order to mitigate this issue, this DLAMI includes functionality that dynamically loads the proprietary driver for G4Dn and G5 instances. Please allow a brief initialization period for this loading in order to ensure that your instances are able to work properly.
 - To check the status and health of this service, you can use the following commands:

```
sudo systemctl is-active dynamic_driver_load.service
active
```

Version 2.17.0

Release date: 2024-09-25

AMI name: Deep Learning OSS Nvidia Driver AMI GPU TensorFlow 2.17 (Ubuntu 22.04) 20240924

Added

- Initial release of the Deep Learning OSS Nvidia Driver AMI GPU TensorFlow 2.17 (Ubuntu 22.04) series.
 - Software Includes the Following:
 - "nvidia-driver=550.90.07"
 - "fabric-manager=550.90.07"
 - "cuda=12.3"
 - "cudnn=8.9.7"
 - "efa=1.34.0"
 - "nccl=2.22.3"
 - "aws-nccl-ofi-plugin=v1.11.0-aws"

- Tensorflow virtual environment (activation command source /opt/tensorflow/bin/activate) includes the following:
 - "tensorflow=2.17.0"

Release Notes for Multi-Framework DLAMIs

🚺 Tip

If you use only one machine learning framework, then we recommend a <u>single-framework</u> DLAMI.

Multi Framework DLAMI Release Notes

Multi-Framework DLAMI Release Notes

Multi-Framework DLAMI Release Notes

Below are the release notes for Multi-Framework X86 DLAMI:

GPU

• Amazon Deep Learning AMI (Amazon Linux 2)

Amazon Neuron

• Refer to the Neuron DLAMI User Guide

Amazon Deep Learning AMI (Amazon Linux 2)

🚺 Tip

Customers using a single framework like PyTorch or TensorFlow are encouraged to use the single framework DLAMIs mentioned <u>here</u>

For help getting started, see Getting started with DLAMI.

AMI name format

- Deep Learning Proprietary Nvidia Driver AMI (Amazon Linux 2) Version \${XX.X}
- Deep Learning OSS Nvidia Driver AMI (Amazon Linux 2) Version \${XX.X}

Supported EC2 instances

- Please refer to Important changes to DLAMI.
- Deep Learning with OSS Nvidia Driver supports G4dn, G5, G6, Gr6, G6e, P4d, P4de, P5
- Deep Learning with Proprietary Nvidia Driver supports G3 (G3.16x not supported), P3, P3dn

- Supported Amazon Service: Amazon EC2
- Operating System: Amazon Linux 2
- Compute Architecture: x86
- Conda environments framework and python versions:
 - Deep Learning OSS Nvidia Driver AMI (Amazon Linux 2):
 - python3: Python 3.10
 - tensorflow2_p310: TensorFlow 2.16, Python 3.10
 - pytorch_p310: PyTorch 2.2, Python 3.10
 - Deep Learning Proprietary Nvidia Driver AMI (Amazon Linux 2):
 - python3: Python 3.10
 - tensorflow2_p310: TensorFlow 2.16, Python 3.10
 - pytorch_p310: PyTorch 2.2, Python 3.10
- NVIDIA Driver:
 - OSS Nvidia driver: 550.163.01
 - Proprietary Nvidia driver: 550.163.01
- NVIDIA CUDA12.1-12.4 stack:
 - CUDA, NCCL and cuDDN installation path: /usr/local/cuda-xx.x/
 - Default CUDA: 12.1
- PATH /usr/local/cuda points to CUDA12.1

- Updated below env vars:
 - LD_LIBRARY_PATH to have /usr/local/cuda-12.1/lib:/usr/local/cuda-12.1/lib64:/usr/local/cuda-12.1:/usr/local/cuda-12.1/targets/x86_64-linux/lib
 - PATH to have /usr/local/cuda-12.1/bin/:/usr/local/cuda-11.8/include/
- For any different CUDA version, please update LD_LIBRARY_PATH accordingly.
- Compiled NCCL Version for CUDA 12.1-12.4: 2.22.3
- NCCL Tests Location:
 - all_reduce, all_gather and reduce_scatter: /usr/local/cuda-xx.x/efa/test-cuda-xx.x/
 - To run NCCL tests, LD_LIBRARY_PATH needs to passed having below updates.
 - Common PATHs are already added to LD_LIBRARY_PATH:
 - /opt/amazon/efa/lib:/opt/amazon/openmpi/lib:/opt/aws-ofi-nccl/lib:/usr/local/lib:/ usr/lib
 - For any different CUDA version, please update LD_LIBRARY_PATH accordingly.
- EFA Installer: 1.38.0
- **GDRCopy:** 2.4
- Amazon OFI NCCL: 1.13.2
 - System location: /usr/local/cuda-xx.x/efa
 - This is added to run NCCL tests located at /usr/local/cuda-xx.x/efa/test-cuda-xx.x/
 - Also, PyTorch package comes with dynamically linked Amazon OFI NCCL plugin as a conda package aws-ofi-nccl-dlc package as well and PyTorch will use that package instead of system Amazon OFI NCCL.
- NCCL Tests Location: /usr/local/cuda-xx.x/efa/test-cuda-xx.x/
- Amazon CLI v2 at /usr/local/bin/aws2 and Amazon CLI v1 at /usr/local/bin/aws
- EBS volume type: gp3
- Query AMI-ID with SSM Parameter (example region is us-east-1):
 - OSS Nvidia Driver:

```
aws ssm get-parameter --name /aws/service/deeplearning/ami/x86_64/multi-framework-
oss-nvidia-driver-amazon-linux-2/latest/ami-id --region us-east-1 --query
"Parameter.Value" --output text
```

```
aws ssm get-parameter --name /aws/service/deeplearning/ami/x86_64/multi-framework-
proprietary-nvidia-driver-amazon-linux-2/latest/ami-id --region us-east-1 --query
"Parameter.Value" --output text
```

- Query AMI-ID with AWSCLI (example region is us-east-1):
 - OSS Nvidia Driver:

```
aws ec2 describe-images --region us-east-1 --owners amazon --filters
'Name=name,Values=Deep Learning OSS Nvidia Driver AMI (Amazon Linux 2)
Version ??.?' 'Name=state,Values=available' --query 'reverse(sort_by(Images, &CreationDate))[:1].ImageId' --output text
```

• Proprietary Nvidia Driver:

```
aws ec2 describe-images --region us-east-1 --owners amazon --filters
'Name=name,Values=Deep Learning Proprietary Nvidia Driver AMI (Amazon Linux 2)
Version ??.?' 'Name=state,Values=available' --query 'reverse(sort_by(Images,
&CreationDate))[:1].ImageId' --output text
```

Notices

EFA Updates from 1.37 to 1.38 (Release on 2025-02-05)

 EFA now bundles the Amazon OFI NCCL plugin, which can now be found in /opt/amazon/ofi-nccl rather than the original /opt/aws-ofi-nccl/. If updating your LD_LIBRARY_PATH variable, please ensure that you modify your OFI NCCL location properly.

Neuron Conda Environment Removal

 Deep Learning Proprietary Nvidia Driver AMIs released after July 18, 2024 will be shipped without neuron conda environments for PyTorch and TensorFlow. Please use the Neuron DLAMIs on the DLAMI Release Notes instead, to utilize neuron environments.

Audit Package Removal

• DLAMI's released between March 26,2024 (2024-03-26) and April 12, 2024 (2024-04-12) were shipped without the audit package. If you require this specific package for your logging and

monitoring needs, please migrate your workflows to the latest DLAMI in order to consume those with the audit package installed.

Horovod

 Horovod is removed from the current pytorch_p310 and tensorflow2_p310 conda environments on the DLAMI. Customers will be able install the horovod libraries by following the <u>horovod</u> <u>guidelines</u> and install them on their DLAMIs for their distributed training jobs.

Release Date: 2025-04-22

AMI names

- Deep Learning OSS Nvidia Driver AMI (Amazon Linux 2) Version 81.2
- Deep Learning Proprietary Nvidia Driver AMI (Amazon Linux 2) Version 81.2

Updated

 Upgraded Nvidia driver from version 550.144.03 to 550.163.01 to address CVEs present in the NVIDIA GPU Display Driver Security Bulletin for April 2025

Release Date: 2025-02-17

AMI names

- Deep Learning OSS Nvidia Driver AMI (Amazon Linux 2) Version 80.6
- Deep Learning Proprietary Nvidia Driver AMI (Amazon Linux 2) Version 80.4

Updated

- Updated NVIDIA Container Toolkit from version 1.17.3 to version 1.17.4
 - Please see the release notes page here for more information: <u>https://github.com/NVIDIA/</u> <u>nvidia-container-toolkit/releases/tag/v1.17.4</u>
 - In Container Toolkit version 1.17.4, the mounting of CUDA compat libraries is now disabled.
 In order to ensure compatibility with multiple CUDA versions on container workflows,
 please ensure you update your LD_LIBRARY_PATH to include your CUDA compatibility
 libraries as shown in under the "If you use a CUDA compatibility layer" tutorial here https://

docs.aws.amazon.com/sagemaker/latest/dg/inference-gpu-drivers.html#collapsible-cudacompat

Removed

 Removed user space libraries cuobj and nvdisasm provided by <u>NVIDIA CUDA toolkit</u> to address CVEs present in the <u>NVIDIA CUDA Toolkit Security Bulletin for February 18, 2025</u>

Release Date: 2025-02-05

AMI names

- Deep Learning Proprietary Nvidia Driver AMI (Amazon Linux 2) Version 80.2
- Deep Learning OSS Nvidia Driver AMI (Amazon Linux 2) Version 80.4

Updated

- Upgraded EFA version from 1.37.0 to 1.38.0
 - EFA now bundles the Amazon OFI NCCL plugin, which can now be found in /opt/amazon/ofinccl rather than the original /opt/aws-ofi-nccl/. If updating your LD_LIBRARY_PATH variable, please ensure that you modify your OFI NCCL location properly.

Release Date: 2025-01-15

AMI names

- Deep Learning OSS Nvidia Driver AMI (Amazon Linux 2) Version 80.3
- Deep Learning Proprietary Nvidia Driver AMI (Amazon Linux 2) Version 80.1

Updated

 Upgraded Nvidia driver from version 550.127.05 to 550.144.03 to address CVEs present in the NVIDIA GPU Display Driver Security Bulletin for January 2025

Release Date: 2024-12-09

AMI names

- Deep Learning OSS Nvidia Driver AMI (Amazon Linux 2) Version 80.1
- Deep Learning Proprietary Nvidia Driver AMI (Amazon Linux 2) Version 79.9

Updated

• Upgraded Nvidia Container Toolkit from version 1.17.0 to 1.17.3

Release Date: 2024-11-11

AMI names

- Deep Learning OSS Nvidia Driver AMI (Amazon Linux 2) Version 79.9
- Deep Learning Proprietary Nvidia Driver AMI (Amazon Linux 2) Version 79.7

Updated

 Upgraded Nvidia Container Toolkit from version 1.16.2 to 1.17.0, addressing the security vulnerability <u>CVE-2024-0134</u>.

Release Date: 2024-10-22

AMI names

- Deep Learning OSS Nvidia Driver AMI (Amazon Linux 2) Version 79.6
- Deep Learning Proprietary Nvidia Driver AMI (Amazon Linux 2) Version 79.6

Updated

 Upgraded Nvidia driver from version 550.90.07 to 550.127.05 to address CVEs present in the NVIDIA GPU Display Security Bulletin for October 2024
Release Date: 2024-10-03

AMI names

- Deep Learning OSS Nvidia Driver AMI (Amazon Linux 2) Version 79.3
- Deep Learning Proprietary Nvidia Driver AMI (Amazon Linux 2) Version 79.3

Updated

• Upgraded Nvidia Container Toolkit from version 1.16.1 to 1.16.2, addressing the security vulnerability CVE-2024-0133.

Release Date: 2024-07-18

AMI names

- Deep Learning OSS Nvidia Driver AMI (Amazon Linux 2) Version 78.6
- Deep Learning Proprietary Nvidia Driver AMI (Amazon Linux 2) Version 78.7

Updated

- Removed aws_neuron_pytorch_p38 and aws_neuron_tensorflow_p38 conda environments from the Deep Learning Proprietary Nvidia Driver AMI.
- Removed Inf1 instance family support from the Deep Learning Proprietary Nvidia Driver AMI.

Release Date: 2024-06-06

AMI names

- Deep Learning OSS Nvidia Driver AMI (Amazon Linux 2) Version 78.5
- Deep Learning Proprietary Nvidia Driver AMI (Amazon Linux 2) Version 78.5

Updated

• Updated Nvidia driver version to 535.183.01 from 535.161.08

Release Date: 2024-05-17

AMI names

- Deep Learning OSS Nvidia Driver AMI (Amazon Linux 2) Version 78.1
- Deep Learning Proprietary Nvidia Driver AMI (Amazon Linux 2) Version 78.1

Updated

• Updated torchserve from v0.8.2 to v0.11.0 in the pytorch_p310 environment.

Release Date: 2024-05-07

AMI names

- Deep Learning OSS Nvidia Driver AMI (Amazon Linux 2) Version 78.0
- Deep Learning Proprietary Nvidia Driver AMI (Amazon Linux 2) Version 78.0

Updated

- TensorFlow version updated from 2.15 to 2.16 in the tensorflow2_p310 environment.
- Updated EFA version from version 1.30 to version 1.32
- Updated Amazon OFI NCCL plugin from version 1.7.4 to version 1.9.1
- Updated Nvidia container toolkit from version 1.13.5 to version 1.15.0
 - NOTE: Version 1.15.0 does NOT include the nvidia-container-runtime and nvidia-docker2 packages. It is recommended to use nvidia-container-toolkit packages directly by following <u>Nvidia container toolkit docs</u>.

Added

• Added CUDA12.3 stack with CUDA12.3, NCCL 2.21.5, CuDNN 8.9.7

Removed

• Removed CUDA11.7, CUDA12.0 stacks present at /usr/local/cuda-11.7 and /usr/local/cuda-12.0

 Removed nvidia-docker2 package and its command nvidia-docker as part of Nvidia container toolkit update from 1.13.5 to <u>1.15.0</u> which does **NOT** include the nvidia-container-runtime and nvidia-docker2 packages.

Release Date: 2024-04-04

AMI names

- Deep Learning OSS Nvidia Driver AMI (Amazon Linux 2) Version 77.0
- Deep Learning Proprietary Nvidia Driver AMI (Amazon Linux 2) Version 77.0

Updated

- PyTorch version updated from 2.1 to 2.2 in the pytorch_p310 environment.
- For OSS Nvidia driver DLAMIs, added G6 and Gr6 EC2 instances support. Please refer EC2 instance selection page for more information.

Release Date: 2024-03-29

AMI names

- Deep Learning OSS Nvidia Driver AMI (Amazon Linux 2) Version 76.8
- Deep Learning Proprietary Nvidia Driver AMI (Amazon Linux 2) Version 76.9

Updated

- Updated Nvidia driver from 535.104.12 to 535.161.08 in both Proprietary and OSS Nvidia driver DLAMIs.
- The new supported instances for each DLAMI are as follows:
 - Deep Learning with Proprietary Nvidia Driver supports G3 (G3.16x not supported), P3, P3dn, Inf1
 - Deep Learning with OSS Nvidia Driver supports G4dn, G5, P4d, P4de.

Removed

• Removed G4dn, G5, G3.16x EC2 instances support from Proprietary Nvidia driver DLAMI.

Version 76.8

Release Date: 2024-03-20

AMI names

• Deep Learning Proprietary Nvidia Driver AMI (Amazon Linux 2) Version 76.8

Added

 Added <u>awscliv2</u> in the AMI as /usr/local/bin/aws2, alongside awscliv1 as /usr/local/bin/aws on Proprietary Nvidia Driver AMI

Version 76.7

Release Date: 2024-03-20

AMI names

• Deep Learning OSS Nvidia Driver AMI (Amazon Linux 2) Version 76.7

Added

- Added <u>awscliv2</u> in the AMI as /usr/local/bin/aws2, alongside awscliv1 as /usr/local/bin/aws on OSS Nvidia Driver AMI
- Updated OSS Nvidia driver DLAMI with G4dn and G5 support, based on it current support looks like below:
 - Deep Learning Base Proprietary Nvidia Driver AMI (Amazon Linux 2) supports P3, P3dn, G3, G5, G4dn.
 - Deep Learning Base OSS Nvidia Driver AMI (Amazon Linux 2) supports G4dn, G5, P4, P5.
- OSS Nvidia driver DLAMIs are recommended to be used for G4dn, G5, P4, P5.

Version 76.3

Release Date: 2024-02-14

Updated

• Updated TensorFlow from 2.13.0 to 2.15.0

- Updated EFA from 1.29.0 to 1.30.0
- Updated Amazon-OFI-NCCL from 1.7.3-aws to 1.7.4-aws
- Updated Nvidia Driver to 535.104.12 on Deep Learning Proprietary Nvidia Driver AMI
- Updated Nvidia Driver to 535.154.05 on Deep Learning OSS Nvidia Driver AMI

Version 76.2

Release Date: 2024-02-02

AMI names

- Deep Learning Proprietary Nvidia Driver AMI (Amazon Linux 2) Version 76.2
- Deep Learning OSS Nvidia Driver AMI (Amazon Linux 2) Version 76.4

Security

• Updated runc package version to consume patch for CVE-2024-21626.

Version 76.1

Release Date: 2023-12-27

Updated

• Updated PyTorch from 2.0.1 to 2.1.0

Version 75.1

Release Date: 2023-11-17

Please refer to Important changes to DLAMI

AMI names

- Deep Learning OSS Nvidia Driver AMI (Amazon Linux 2) Version 75.1
- Deep Learning Proprietary Nvidia Driver AMI (Amazon Linux 2) Version 75.1

Added

- Amazon Deep Learning AMI (DLAMI) is split into two separate groups:
 - DLAMI that uses Nvidia Proprietary Driver (to support P3, P3dn, G3, G5, G4dn).
 - DLAMI that uses Nvidia OSS Driver to enable EFA (to support P4, P5).
- Please refer to public announcement for more information on DLAMI split.
- Amazon cli queries for above are in the <u>release notes</u> under bullet point Query AMI-ID with AWSCLI (example region is us-east-1)

Updated

- EFA updated from 1.26.1 to 1.29.0
- GDRCopy updated from 2.3 to 2.4

Version 74.4

Release Date: 2023-10-27

Updated

- Amazon OFI NCCL Plugin updated from version 1.7.2 to version 1.7.3
- Updated CUDA 12.0-12.1 directories with NCCL version 2.18.5
- CUDA12.1 updated as the default CUDA Version
 - Updated LD_LIBRARY_PATH to have /usr/local/cuda-12.1/targets/x86_64-linux/lib/:/usr/local/cuda-12.1/lib:/usr/local/cuda-12.1/lib64:/usr/local/cuda-12.1 and PATH to have /usr/local/cuda-12.1/bin/
 - For customers looking to change to any different CUDA version, please define the LD_LIBRARY_PATH and PATH variables accordingly.
- Updated Pillow from version 9.4.0 to 10.1.0 to fix <u>SNYK-PYTHON-PILLOW-5918878</u> in all conda environments
- Updated opencv-python from 4.8.0.74 to 4.8.1.78 to fix <u>SNYK-PYTHON-OPENCVPYTHON-5926695</u> in all conda environments

Added

- Kernel Live Patching is now enabled. Live patching enables customers to apply security vulnerability and critical bug patches to a running Linux kernel, without reboots or disruptions to running applications.
 - Please note that live patching support for kernel 5.10.192 will end on 11/30/23.
 - For more information please reference the official Amazon documents here <u>https://</u> docs.aws.amazon.com/AWSEC2/latest/UserGuide/al2-live-patching.html

Version 74.0

Release Date: 2023-07-19

Updated

- Updated TensorFlow from 2.12 to 2.13
 - Horovod has been removed from the conda environment in this release. See **Notice** for details on installing horovod.

Version 73.1

Release Date: 2023-06-12

Updated

• Updated PyTorch from 2.0.0 to 2.0.1

Getting started with DLAMI

This guide includes tips about picking the DLAMI that's right for you, selecting an instance type that fits your use case and budget, and <u>Related information about DLAMI</u> that describes custom setups that may be of interest.

If you're new to using Amazon or using Amazon EC2, start with the <u>Deep Learning AMI with Conda</u>. If you're familiar with Amazon EC2 and other Amazon services like Amazon EMR, Amazon EFS, or Amazon S3, and are interested in integrating those services for projects that need distributed training or inference, then check out <u>Related information about DLAMI</u> to see if one fits your use case.

We recommend that you check out <u>Choosing a DLAMI</u> to get an idea of which instance type might be best for your application.

Next step

Choosing a DLAMI

Choosing a DLAMI

We offer a range of DLAMI options as mentioned in the <u>GPU DLAMI release notes</u>. To help you select the correct DLAMI for your use case, we group images by the hardware type or functionality for which they were developed. Our top level groupings are:

- DLAMI Type: Base, Single-Framework, Multi-Framework (Conda DLAMI)
- Compute Architecture: x86-based, Arm64-based Amazon Graviton
- Processor Type: <u>GPU</u>, <u>CPU</u>, <u>Inferentia</u>, <u>Trainium</u>
- SDK: CUDA, Amazon Neuron
- **OS:** Amazon Linux, Ubuntu

The rest of the topics in this guide help further inform you and go into more details.

Topics

- CUDA Installations and Framework Bindings
- Deep Learning Base AMI

- Deep Learning AMI with Conda
- DLAMI Architecture Options
- DLAMI Operating System Options

Next Up

Deep Learning AMI with Conda

CUDA Installations and Framework Bindings

While deep learning is all pretty cutting edge, each framework offers "stable" versions. These stable versions may not work with the latest CUDA or cuDNN implementation and features. Your use case and the features you require can help you choose a framework. If you are not sure, then use the latest Deep Learning AMI with Conda. It has official pip binaries for all frameworks with CUDA, using whichever most recent version is supported by each framework. If you want the latest versions, and to customize your deep learning environment, use the Deep Learning Base AMI.

Look at our guide on <u>Stable Versus Release Candidates</u> for further guidance.

Choose a DLAMI with CUDA

The Deep Learning Base AMI has all available CUDA version series

The Deep Learning AMI with Conda has all available CUDA version series

🚺 Note

We no longer include the MXNet, CNTK, Caffe, Caffe2, Theano, Chainer, or Keras Conda environments in the Amazon Deep Learning AMIs.

For specific framework version numbers, see the Deep Learning AMIs Release Notes

Choose this DLAMI type or learn more about the different DLAMIs with the **Next Up** option.

Choose one of the CUDA versions and review the full list of DLAMIs that have that version in the **Appendix**, or learn more about the different DLAMIs with the **Next Up** option.

Next Up

CUDA Installations and Framework Bindings

Deep Learning Base AMI

Related Topics

 For instructions on switching between CUDA versions, refer to the <u>Using the Deep Learning Base</u> <u>AMI</u> tutorial.

Deep Learning Base AMI

The Deep Learning Base AMI is like an empty canvas for deep learning. It comes with everything you need up until the point of the installation of a particular framework, and has your choice of CUDA versions.

Why to Choose the Base DLAMI

This AMI group is useful for project contributors who want to fork a deep learning project and build the latest. It's for someone who wants to roll their own environment with the confidence that the latest NVIDIA software is installed and working so they can focus on picking which frameworks and versions they want to install.

Choose this DLAMI type or learn more about the different DLAMIs with the **Next Up** option.

Next Up

DLAMI with Conda

Related Topics

Using the Deep Learning Base AMI

Deep Learning AMI with Conda

The Conda DLAMI uses conda virtual environments, they are present either multi-framework or single framework DLAMIs. These environments are configured to keep the different framework installations separate and streamline switching between frameworks. This is great for learning and experimenting with all of the frameworks the DLAMI has to offer. Most users find that the new Deep Learning AMI with Conda is perfect for them.

They are updated often with the latest versions from the frameworks, and have the latest GPU drivers and software. They are generally referred to as <u>the</u> Amazon Deep Learning AMIs in most

documents. These DLAMIs support Ubuntu 20.04, Ubuntu 22.04, Amazon Linux 2, Amazon Linux 2023 Operating systems. Operating systems support depends on support from upstream OS.

Stable Versus Release Candidates

The Conda AMIs use optimized binaries of the most recent formal releases from each framework. Release candidates and experimental features are not to be expected. The optimizations depend on the framework's support for acceleration technologies like Intel's MKL DNN, which speeds up training and inference on C5 and C4 CPU instance types. The binaries are also compiled to support advanced Intel instruction sets including but not limited to AVX, AVX-2, SSE4.1, and SSE4.2. These accelerate vector and floating point operations on Intel CPU architectures. Additionally, for GPU instance types, the CUDA and cuDNN are updated with whichever version the latest official release supports.

The Deep Learning AMI with Conda automatically installs the most optimized version of the framework for your Amazon EC2 instance upon the framework's first activation. For more information, refer to Using the Deep Learning AMI with Conda.

If you want to install from source, using custom or optimized build options, the <u>Deep Learning Base</u> <u>AMIs</u> might be a better option for you.

Python 2 Deprecation

The Python open source community has officially ended support for Python 2 on January 1, 2020. The TensorFlow and PyTorch community have announced that the TensorFlow 2.1 and PyTorch 1.4 releases are the last ones supporting Python 2. Previous releases of the DLAMI (v26, v25, etc) that contain Python 2 Conda environments continue to be available. However, we provide updates to the Python 2 Conda environments on previously published DLAMI versions only if there are security fixes published by the open-source community for those versions. DLAMI releases with the latest versions of the TensorFlow and PyTorch frameworks do not contain the Python 2 Conda environments.

CUDA Support

Specific CUDA version numbers can be found in the <u>GPU DLAMI release notes</u>.

Next Up

DLAMI Architecture Options

Related Topics

 For a tutorial on using a Deep Learning AMI with Conda, see the <u>Using the Deep Learning AMI</u> with Conda tutorial.

DLAMI Architecture Options

Amazon Deep Learning AMIss are offered with either x86-based or Arm64-based <u>Amazon</u> <u>Graviton2</u> architectures.

For information about getting started with the ARM64 GPU DLAMI, see <u>The ARM64 DLAMI</u>. For more details on available instance types, see <u>Choosing a DLAMI instance type</u>.

Next Up

DLAMI Operating System Options

DLAMI Operating System Options

DLAMIs are offered in the following operating systems.

- Amazon Linux 2
- Amazon Linux 2023
- Ubuntu 20.04
- Ubuntu 22.04

Older versions of operating systems are available on deprecated DLAMIs. For more information on DLAMI deprecation, see Deprecations for DLAMI

Before choosing a DLAMI, assess what instance type you need and identify your Amazon Region.

Next Up

Choosing a DLAMI instance type

Choosing a DLAMI instance type

More generally, consider the following when choosing an instance type for a DLAMI.

- If you're new to deep learning, then an instance with a single GPU might suit your needs.
- If you're budget conscious, then you can use CPU-only instances.
- If you're looking to optimize high performance and cost efficiency for deep learning model inference, then you can use instances with Amazon Inferentia chips.
- If you're looking for a high performance GPU instance with an Arm64-based CPU architecture, then you can use the G5g instance type.
- For high-volume inference services, a single CPU instance with a lot of memory, or a cluster of such instances, might be a better solution.
- If you're using a large model with a lot of data or a high batch size, then you need a larger instance with more memory. You can also distribute your model to a cluster of GPUs. You may find that using an instance with less memory is a better solution for you if you decrease your batch size. This may impact your accuracy and training speed.
- If you're interested in running machine learning applications using NVIDIA Collective Communications Library (NCCL) requiring high levels of inter-node communications at scale, you might want to use <u>Elastic Fabric Adapter (EFA)</u>.

For more detail on instances, see EC2 Instance Types.

The following topics provide information about instance type considerations.

<u> Important</u>

The Deep Learning AMIs include drivers, software, or toolkits developed, owned, or provided by NVIDIA Corporation. You agree to use these NVIDIA drivers, software, or toolkits only on Amazon EC2 instances that include NVIDIA hardware.

Topics

- Pricing for the DLAMI
- DLAMI Region Availability
- <u>Recommended GPU Instances</u>
- <u>Recommended CPU Instances</u>
- <u>Recommended Inferentia Instances</u>
- Recommended Trainium Instances

Pricing for the DLAMI

The deep learning frameworks included in the DLAMI are free, and each has its own opensource licenses. Although the software included in the DLAMI is free, you still have to pay for the underlying Amazon EC2 instance hardware.

Some Amazon EC2 instance types are labeled as free. It is possible to run the DLAMI on one of these free instances. This means that using the DLAMI is entirely free when you only use that instance's capacity. If you need a more powerful instance with more CPU cores, more disk space, more RAM, or one or more GPUs, then you need an instance that is not in the free-tier instance class.

For more information about instance choice and pricing, see <u>Amazon EC2 pricing</u>.

DLAMI Region Availability

Each Region supports a different range of instance types and often an instance type has a slightly different cost in different Regions. DLAMIs are not available in every Region, but it is possible to copy DLAMIs to the Region of your choice. See <u>Copying an AMI</u> for more information. Note the Region selection list and be sure you pick a Region that's close to you or your customers. If you plan to use more than one DLAMI and potentially create a cluster, be sure to use the same Region for all of nodes in the cluster.

For a more info on Regions, visit EC2 Regions.

Next Up

Recommended GPU Instances

Recommended GPU Instances

We recommend a GPU instance for most deep learning purposes. Training new models is faster on a GPU instance than a CPU instance. You can scale sub-linearly when you have multi-GPU instances or if you use distributed training across many instances with GPUs.

The following instance types support the DLAMI. For information about GPU instance type options and their uses, see <u>EC2 Instance Types</u> and select **Accelerated Computing**.

🚯 Note

The size of your model should be a factor in choosing an instance. If your model exceeds an instance's available RAM, choose a different instance type with enough memory for your application.

DLAMI instances provide tooling to monitor and optimize your GPU processes. For more information about monitoring your GPU processes, see <u>GPU Monitoring and Optimization</u>.

For specific tutorials on working with G5g instances, see <u>The ARM64 DLAMI</u>.

Next Up

Recommended CPU Instances

Recommended CPU Instances

Whether you're on a budget, learning about deep learning, or just want to run a prediction service, you have many affordable options in the CPU category. Some frameworks take advantage of Intel's MKL DNN, which speeds up training and inference on C5 (not available in all Regions) CPU instance types. For information about CPU instance types, see <u>EC2 Instance Types</u> and select **Compute Optimized**.

🚺 Note

The size of your model should be a factor in choosing an instance. If your model exceeds an instance's available RAM, choose a different instance type with enough memory for your application.

Next Up

Recommended Inferentia Instances

Recommended Inferentia Instances

Amazon Inferentia instances are designed to provide high performance and cost efficiency for deep learning model inference workloads. Specifically, Inf2 instance types use Amazon Inferentia chips and the <u>Amazon Neuron SDK</u>, which is integrated with popular machine learning frameworks such as TensorFlow and PyTorch.

Customers can use Inf2 instances to run large scale machine learning inference applications such as search, recommendation engines, computer vision, speech recognition, natural language processing, personalization, and fraud detection, at the lowest cost in the cloud.

1 Note

The size of your model should be a factor in choosing an instance. If your model exceeds an instance's available RAM, choose a different instance type with enough memory for your application.

 <u>Amazon EC2 Inf2 Instances</u> have up to up to 16 Amazon Inferentia chips and 100 Gbps of networking throughput.

For more information about getting started with Amazon Inferentia DLAMIs, see <u>The Amazon</u> Inferentia Chip With DLAMI.

Next Up

Recommended Trainium Instances

Recommended Trainium Instances

Amazon Trainium instances are designed to provide high performance and cost efficiency for deep learning model inference workloads. Specifically, Trn1 instance types use Amazon Trainium chips and the <u>Amazon Neuron SDK</u>, which is integrated with popular machine learning frameworks such as TensorFlow and PyTorch.

Customers can use Trn1 instances to run large scale machine learning inference applications such as search, recommendation engines, computer vision, speech recognition, natural language processing, personalization, and fraud detection, at the lowest cost in the cloud.

(i) Note

The size of your model should be a factor in choosing an instance. If your model exceeds an instance's available RAM, choose a different instance type with enough memory for your application.

 <u>Amazon EC2 Trn1 Instances</u> have up to up to 16 Amazon Trainium chips and 100 Gbps of networking throughput.

Setting up a DLAMI instance

After you <u>choose a DLAMI</u> and <u>choose an Amazon Elastic Compute Cloud (Amazon EC2) instance</u> <u>type</u> that you want to use, then you're ready to set up your new DLAMI instance.

If you haven't yet chosen a DLAMI and EC2 instance type, then see Getting started with DLAMI.

Topics

- Finding the ID of a DLAMI
- Launching a DLAMI instance
- Connecting to a DLAMI instance
- Setting up a Jupyter Notebook server on a DLAMI instance
- Cleaning up a DLAMI instance

Finding the ID of a DLAMI

Each DLAMI has a unique identifier (ID). When you launch a DLAMI instance using the Amazon EC2 console, you can optionally use the DLAMI ID to search for the DLAMI that you want to use. When you launch a DLAMI instance using the Amazon Command Line Interface (Amazon CLI), this ID is required.

You can find the ID for the DLAMI of your choice by using an Amazon CLI command for either Amazon EC2 or Parameter Store, a capability of Amazon Systems Manager. For instructions on installing and configuring the Amazon CLI, see <u>Get started with the Amazon CLI</u> in the *Amazon Command Line Interface User Guide*.

Using Parameter Store

To find a DLAMI ID using ssm get-parameter

In the following <u>ssm get-parameter</u> command, for the --name option, the parameter name format is /aws/service/deeplearning/ami/\$architecture/\$ami_type/latest/ ami-id. In this name format, architecture can be either **x86_64** or **arm64**. Specify the ami_type by taking the DLAMI name and removing the keywords "deep", "learning", and "ami". AMI Name can be found in Deep Learning AMIs Release Notes.

A Important

To use this command, the Amazon Identity and Access Management (IAM) principal that you use must have the ssm:GetParameter permission. For more information about IAM principals, see the Additional resources section of **IAM roles** in the *IAM User Guide*.

```
aws ssm get-parameter --name /aws/service/deeplearning/ami/x86_64/base-oss-
nvidia-driver-ubuntu-22.04/latest/ami-id \
--region us-east-1 --query "Parameter.Value" --output text
```

The output should look similar to the following:

ami-09ee1a996ac214ce7

🚺 Tip

For some currently supported DLAMI frameworks, you can find more specific example **ssm get-parameter** commands in <u>Deep Learning AMIs Release Notes</u>. Choose the link to the release notes of your chosen DLAMI, and then find its ID query in the release notes.

Using Amazon EC2 CLI

To find a DLAMI ID using ec2 describe-images

In the following <u>ec2 describe-images</u> command, for the value of the filter Name=name, enter the DLAMI name. You can specify a release version for a given framework, or you can get the latest release by replacing the version number with a question mark (?).

```
•
```

```
aws ec2 describe-images --region us-east-1 --owners amazon \
--filters 'Name=name,Values=Deep Learning Base OSS Nvidia Driver GPU AMI (Ubuntu
22.04) ???????' 'Name=state,Values=available' \
--query 'reverse(sort_by(Images, &CreationDate))[:1].ImageId' --output text
```

The output should look similar to the following:

Developer Guide

ami-09ee1a996ac214ce7

🚺 Tip

For an example **ec2 describe-images** command that's specific to the DLAMI of your choice, see <u>Deep Learning AMIs Release Notes</u>. Choose the link to the release notes of your chosen DLAMI, and then find its ID query in the release notes.

Next step

Launching a DLAMI instance

Launching a DLAMI instance

After you <u>find the ID</u> of the DLAMI that you want to use to launch a DLAMI instance, you're ready to launch the instance. To launch it, you can use either the Amazon EC2 console or the Amazon Command Line Interface (Amazon CLI).

Note

For this walkthrough, we might make references specific to the Deep Learning Base OSS Nvidia Driver GPU AMI (Ubuntu 22.04). Even if you select a different DLAMI, you should be able to follow this guide.

EC2 console

Note

To accelerate high-performance computing (HPC) and machine learning applications, you can launch your DLAMI instance with an Elastic Fabric Adapter (EFA). For specific instructions, see Launching a Amazon Deep Learning AMIs Instance With EFA.

1. Open the EC2 Console.

- 2. Note your current Amazon Web Services Region in the top-most navigation. If this isn't your desired Region, then change this option before you continue. For more information, see EC2 Regions in the Amazon Web Services General Reference.
- 3. Choose Launch Instance.
- 4. Enter a name for your instance and select the DLAMI that is right for you.
 - a. Find an existing DLAMI in **My AMIs** or choose **Quick Start.**
 - b. Search by DLAMI ID. Browse the options then select your choice.
- Choose an instance type. You can find the recommended instance families for your DLAMI in <u>Deep Learning AMIs Release Notes</u>. For general recommendations on DLAMI instance types, see Choosing a DLAMI instance type.
- 6. Choose Launch Instance.

Amazon CLI

 To use the Amazon CLI, you must have the ID of the DLAMI that you want to use, the Amazon Web Services Region and EC2 instance type, and your security token information. Then, you can launch the instance using the ec2 run-instances Amazon CLI command.

For instructions on installing and configuring the Amazon CLI, see <u>Get started with the</u> <u>Amazon CLI</u> in the Amazon Command Line Interface User Guide. For more information, including command examples, see <u>Launch</u>, list, and close Amazon EC2 instances for the Amazon CLI.

After you launch your instance using either the Amazon EC2 console or Amazon CLI, wait for the instance to be ready. This usually takes only a few minutes. You can verify the status of the instance in the <u>Amazon EC2 console</u>. For more information, see <u>Status checks for Amazon EC2</u> instances in the *Amazon EC2 User Guide*.

Next step

Connecting to a DLAMI instance

Connecting to a DLAMI instance

After you <u>launch a DLAMI instance</u> and the instance is running, you can connect to it from a client (Windows, macOS, or Linux) using SSH. For instructions, see <u>Connect to your Linux instance using</u> <u>SSH</u> in the *Amazon EC2 User Guide*.

Keep a copy of the SSH login command handy in case you want to set up a Jupyter Notebook server after you log in. To connect to the Jupyter webpage, you use a variation of that command.

Next step

Setting up a Jupyter Notebook server on a DLAMI instance

Setting up a Jupyter Notebook server on a DLAMI instance

With a Jupyter Notebook server, you can create and run Jupyter notebooks from your DLAMI instance. With Jupyter notebooks, you can conduct machine learning (ML) experiments for training and inference while using the Amazon infrastructure and accessing packages built into the DLAMI. For more information about Jupyter notebooks, see <u>The Jupyter Notebook</u> on the Jupyter User Documentation website.

To set up a Jupyter Notebook server, you must:

- Configure the Jupyter Notebook server on your DLAMI instance.
- Configure your client to connect to the Jupyter Notebook server. We provide configuration instructions for Windows, macOS, and Linux clients.
- Test the setup by logging in to the Jupyter Notebook server.

To complete these steps, follow the instructions in the following topics. After you've set up a Jupyter Notebook server, you can run the example notebook tutorials that ship in the DLAMIs. For more information, see Running Jupyter Notebook Tutorials.

Topics

- Securing the Jupyter Notebook server on a DLAMI instance
- Starting the Jupyter Notebook server on a DLAMI instance
- Connecting a client to the Jupyter Notebook server on a DLAMI instance
- Logging in to the Jupyter Notebook server on a DLAMI instance

Securing the Jupyter Notebook server on a DLAMI instance

To keep your Jupyter Notebook server secure, we recommend setting up a password and creating an SSL certificate for the server. To configure a password and SSL, first <u>connect to your DLAMI</u> <u>instance</u>, and then follow these instructions.

To secure the Jupyter Notebook server

1. Jupyter provides a password utility. Run the following command and enter your preferred password at the prompt.

\$ jupyter notebook password

The output will look something like this:

```
Enter password:
    Verify password:
    [NotebookPasswordApp] Wrote hashed password to /home/ubuntu/.jupyter/
jupyter_notebook_config.json
```

2. Create a self-signed SSL certificate. Follow the prompts to fill out your locality as you see fit. You must enter . if you wish to leave a prompt blank. Your answers will not impact the functionality of the certificate.

```
$ cd ~
    $ mkdir ssl
    $ cd ssl
    $ openssl req -x509 -nodes -days 365 -newkey rsa:2048 -keyout mykey.key -out
    mycert.pem
```

Note

You might be interested in creating a regular SSL certificate that is third-party signed and does not cause the browser to give you a security warning. This process is much more involved. For more information, see <u>Securing a notebook server</u> in the Jupyter Notebook user documentation.

Next step

Starting the Jupyter Notebook server on a DLAMI instance

Starting the Jupyter Notebook server on a DLAMI instance

After you <u>secure your Jupyter Notebook server with a password and SSL</u>, you can start the server. Log in to your DLAMI instance and run the following command that uses the SSL certificate that you created previously.

```
$ jupyter notebook --certfile=~/ssl/mycert.pem --keyfile ~/ssl/mykey.key
```

With the server started, you can now connect to it via an SSH tunnel from your client computer. When the server runs, you will see some output from Jupyter confirming that the server is running. At this point, ignore the callout that you can access the server via a local host URL, because that won't work until you create the tunnel.

🚯 Note

Jupyter will handle switching environments for you when you switch frameworks using the Jupyter web interface. For more information, see <u>Switching Environments with Jupyter</u>.

Next step

Connecting a client to the Jupyter Notebook server on a DLAMI instance

Connecting a client to the Jupyter Notebook server on a DLAMI instance

After you <u>start the Jupyter Notebook server on your DLAMI instance</u>, configure your Windows, macOS, or Linux client to connect to the server. When you connect, you can create and access Jupyter notebooks on the server in your workspace and run your deep learning code on the server.

Prerequisites

Be sure you have the following, which you need to set up an SSH tunnel:

• The public DNS name of your Amazon EC2 instance. For more information, see <u>Amazon EC2</u> instance hostname types in the *Amazon EC2 User Guide*.

• The key pair for the private key file. For more information about accessing your key pair, see Amazon EC2 key pairs and Amazon EC2 instances in the Amazon EC2 User Guide.

Connect from a Windows, macOS, or Linux client

To connect to your DLAMI instance from a Windows, macOS, or Linux client, follow the instructions for your client's operating system.

Windows

To connect to your DLAMI instance from a Windows client using SSH

- Use an SSH client for Windows, such as PuTTY. For instructions, see <u>Connect to your Linux</u> <u>instance using PuTTY</u> in the *Amazon EC2 User Guide*. For other SSH connection options, see <u>Connect to your Linux instance using SSH</u>.
- 2. (Optional) Create an SSH tunnel to a running Jupyter server. Install Git Bash on your Windows client, and then follow the connection instructions for macOS and Linux clients.

macOS or Linux

To connect to your DLAMI instance from a macOS or Linux client using SSH

- 1. Open a terminal.
- 2. Run the following command to forward all requests on local port 8888 to port 8888 on your remote Amazon EC2 instance. Update the command by replacing the location of your key to access the Amazon EC2 instance and the public DNS name of your Amazon EC2 instance. Note, for an Amazon Linux AMI, the user name is ec2-user instead of ubuntu.

```
$ ssh -i ~/mykeypair.pem -N -f -L 8888:localhost:8888 ubuntu@ec2-###-##-
###.compute-1.amazonaws.com
```

This command opens a tunnel between your client and the remote Amazon EC2 instance that is running the Jupyter Notebook server.

Next step

Logging in to the Jupyter Notebook server on a DLAMI instance

Logging in to the Jupyter Notebook server on a DLAMI instance

After you <u>connect your client to the Jupyter Notebook server on your DLAMI instance</u>, you can log in to the server.

To log in to the server in your browser

- 1. In the address bar of your browser, enter the following URL, or click on this link: https://localhost:8888
- 2. With a self-signed SSL certificate, your browser will warn you and prompt you to avoid continuing to visit the website.



Your connection is not private Attackers might be trying to steal your information from localhost (for example, passwords, messages, or credit cards). Learn more NET::ERR_CERT_AUTHORITY_INVALID Help improve Safe Browsing by sending some system information and page content to Google. Privacy policy Advanced

Since you set this up yourself, it is safe to continue. Depending your browser you will get an "advanced", "show details", or similar button.



Your connection is not private

Attackers might be trying to steal your information from **localhost** (for example, passwords, messages, or credit cards). Learn more

NET::ERR_CERT_AUTHORITY_INVALID

Help improve Safe Browsing by sending some <u>system information and page content</u> to Google. <u>Privacy policy</u>

Hide advanced

Back to safety

This server could not prove that it is **localhost**; its security certificate is not trusted by your computer's operating system. This may be caused by a misconfiguration or an attacker intercepting your connection.

Proceed to localhost (unsafe)

Click on this, then click on the "proceed to localhost" link. If the connection is successful, you see the Jupyter Notebook server webpage. At this point, you will be asked for the password you previously set up.

Now you have access to the Jupyter Notebook server that is running on the DLAMI instance. You can create new notebooks or run the provided Tutorials.

Cleaning up a DLAMI instance

When you no longer need your DLAMI instance, you can stop it or terminate it on Amazon EC2 to avoid incurring unexpected charges.

If you stop an instance, you can keep it around and start it later when you want to use it again. Your configurations, files, and other non-volatile information are stored in a volume on Amazon Simple Storage Service (Amazon S3). While your instance is stopped, you incur S3 charges for retaining the volume, but you don't incur charges for compute resources. When your start the instance again, it will mount that storage volume with your data.

If you terminate an instance, it's gone, and you cannot start it again. Of course, you won't incur any more charges for the compute resources with a terminated instance. However, your data still resides on Amazon S3, and you can continue to incur S3 charges. To prevent all further charges related to your terminated instance, you must also delete the storage volume on Amazon S3. For instructions, see <u>Terminate Amazon EC2 instances</u> in the *Amazon EC2 User Guide*.

For more information about Amazon EC2 instance states, such as stopped and terminated, see <u>Amazon EC2 instance state changes</u> in the *Amazon EC2 User Guide*.

Using a DLAMI

Topics

- Using the Deep Learning AMI with Conda
- Using the Deep Learning Base AMI
- <u>Running Jupyter Notebook Tutorials</u>
- Tutorials

The following sections describe how the Deep Learning AMI with Conda can be used to switch environments, run sample code from each of the frameworks, and run Jupyter so you can try out different notebook tutorials.

Using the Deep Learning AMI with Conda

Topics

- Introduction to the Deep Learning AMI with Conda
- Log in to Your DLAMI
- Start the TensorFlow Environment
- Switch to the PyTorch Python 3 Environment
- <u>Removing Environments</u>

Introduction to the Deep Learning AMI with Conda

Conda is an open source package management system and environment management system that runs on Windows, macOS, and Linux. Conda quickly installs, runs, and updates packages and their dependencies. Conda easily creates, saves, loads and switches between environments on your local computer.

The Deep Learning AMI with Conda has been configured for you to easily switch between deep learning environments. The following instructions guide you on some basic commands with conda. They also help you verify that the basic import of the framework is functioning, and that you can run a couple simple operations with the framework. You can then move on to more thorough

tutorials provided with the DLAMI or the frameworks' examples found on each frameworks' project site.

Log in to Your DLAMI

After you log in to your server, you will see a server "message of the day" (MOTD) describing various Conda commands that you can use to switch between the different deep learning frameworks. Below is an example MOTD. Your specific MOTD may vary as new versions of the DLAMI are released.

```
______
      AMI Name: Deep Learning OSS Nvidia Driver AMI (Amazon Linux 2) Version 77
       Supported EC2 instances: G4dn, G5, G6, Gr6, P4d, P4de, P5
          * To activate pre-built tensorflow environment, run: 'source activate
tensorflow2_p310'
          * To activate pre-built pytorch environment, run: 'source activate
pytorch_p310'
          * To activate pre-built python3 environment, run: 'source activate python3'
       NVIDIA driver version: 535.161.08
   CUDA versions available: cuda-11.7 cuda-11.8 cuda-12.0 cuda-12.1 cuda-12.2
   Default CUDA version is 12.1
   Release notes: https://docs.aws.amazon.com/dlami/latest/devguide/appendix-ami-
release-notes.html
   AWS Deep Learning AMI Homepage: https://aws.amazon.com/machine-learning/amis/
   Developer Guide and Release Notes: https://docs.aws.amazon.com/dlami/latest/
devguide/what-is-dlami.html
   Support: https://forums.aws.amazon.com/forum.jspa?forumID=263
   For a fully managed experience, check out Amazon SageMaker at https://
aws.amazon.com/sagemaker
```

Start the TensorFlow Environment

1 Note

When you launch your first Conda environment, please be patient while it loads. The Deep Learning AMI with Conda automatically installs the most optimized version of the

framework for your EC2 instance upon the framework's first activation. You should not expect subsequent delays.

1. Activate the TensorFlow virtual environment for Python 3.

\$ source activate tensorflow2_p310

2. Start the iPython terminal.

(tensorflow2_p310)\$ ipython

3. Run a quick TensorFlow program.

```
import tensorflow as tf
hello = tf.constant('Hello, TensorFlow!')
sess = tf.Session()
print(sess.run(hello))
```

You should see "Hello, Tensorflow!"

Next Up

Running Jupyter Notebook Tutorials

Switch to the PyTorch Python 3 Environment

If you're still in the iPython console, use quit(), then get ready to switch environments.

• Activate the PyTorch virtual environment for Python 3.

```
$ source activate pytorch_p310
```

Test Some PyTorch Code

To test your installation, use Python to write PyTorch code that creates and prints an array.

1. Start the iPython terminal.

(pytorch_p310)\$ ipython

2. Import PyTorch.

```
import torch
```

You might see a warning message about a third-party package. You can ignore it.

3. Create a 5x3 matrix with the elements initialized randomly. Print the array.

```
x = torch.rand(5, 3)
print(x)
```

Verify the result.

tensor([[0.3105, 0.5983, 0.5410], [0.0234, 0.0934, 0.0371], [0.9740, 0.1439, 0.3107], [0.6461, 0.9035, 0.5715], [0.4401, 0.7990, 0.8913]])

Removing Environments

If you run out of space on the DLAMI, you can choose to uninstall Conda packages that you are not using:

```
conda env list
conda env remove --name <env_name>
```

Using the Deep Learning Base AMI

Using the Deep Learning Base AMI

The Base AMI comes with a foundational platform of GPU drivers and acceleration libraries to deploy your own customized deep learning environment. By default the AMI is configured with any one NVIDIA CUDA version environment. You can also switch between different versions of CUDA. Refer to the following instructions for how to do this.

Configuring CUDA Versions

You can verify the CUDA version by running NVIDIA's nvcc program.

```
nvcc --version
```

You can select and verify a particular CUDA version with the following bash command:

```
sudo rm /usr/local/cuda
sudo ln -s /usr/local/cuda-12.0 /usr/local/cuda
```

For more information, see the **Base DLAMI release notes**.

Running Jupyter Notebook Tutorials

Tutorials and examples ship with each of the deep learning projects' source and in most cases they will run on any DLAMI. If you chose the <u>Deep Learning AMI with Conda</u>, you get the added benefit of a few hand-picked tutorials already set up and ready to try out.

<u> Important</u>

To run the Jupyter notebook tutorials installed on the DLAMI, you will need to <u>Setting up a</u> Jupyter Notebook server on a DLAMI instance.

Once the Jupyter server is running, you can run the tutorials through your web browser. If you are running the Deep Learning AMI with Conda or if you have set up Python environments, you can switch Python kernels from the Jupyter notebook interface. Select the appropriate kernel before trying to run a framework-specific tutorial. Further examples of this are provided for users of the Deep Learning AMI with Conda.

Note

Many tutorials require additional Python modules that may not be set up on your DLAMI. If you get an error like "xyz module not found", log in to the DLAMI, activate the environment as described above, then install the necessary modules.

🚺 Tip

Deep learning tutorials and examples often rely on one or more GPUs. If your instance type doesn't have a GPU, you may need to change some of the example's code to get it to run.

Navigating the Installed Tutorials

Once you're logged in to the Jupyter server and can see the tutorials directory (on Deep Learning AMI with Conda only), you will be presented with folders of tutorials by each framework name. If you don't see a framework listed, then tutorials are not available for that framework on your current DLAMI. Click on the name of the framework to see the listed tutorials, then click a tutorial to launch it.

The first time you run a notebook on the Deep Learning AMI with Conda, it will want to know which environment you would like to use. It will prompt you to select from a list. Each environment is named according to this pattern:

Environment (conda_framework_python-version)

For example, you might see Environment (conda_mxnet_p36), which signifies that the environment has MXNet and Python 3. The other variation of this would be Environment (conda_mxnet_p27), which signifies that the environment has MXNet and Python 2.

🚺 Tip

If you're concerned about which version of CUDA is active, one way to see it is in the MOTD when you first log in to the DLAMI.

Switching Environments with Jupyter

If you decide to try a tutorial for a different framework, be sure to verify the currently running kernel. This info can be seen in the upper right of the Jupyter interface, just below the logout button. You can change the kernel on any open notebook by clicking the Jupyter menu item **Kernel**, then **Change Kernel**, and then clicking the environment that suits the notebook you're running.

At this point you'll need to rerun any cells because a change in the kernel will erase the state of anything you've run previously.

🚺 Tip

Switching between frameworks can be fun and educational, however you can run out of memory. If you start running into errors, look at the terminal window that has the Jupyter server running. There are helpful messages and error logging here, and you may see an out-of-memory error. To fix this, you can go to the home page of your Jupyter server, click the **Running** tab, then click **Shutdown** for each of the tutorials that are probably still running in the background and eating up all of your memory.

Tutorials

The following are tutorials on how to use the Deep Learning AMI with Conda's software.

Topics

- <u>Activating Frameworks</u>
- Distributed training using Elastic Fabric Adapter
- GPU Monitoring and Optimization
- The Amazon Inferentia Chip With DLAMI
- The ARM64 DLAMI
- Inference
- Model Serving

Activating Frameworks

The following are the deep learning frameworks installed on the Deep Learning AMI with Conda. Click on a framework to learn how to activate it.

Topics

- PyTorch
- TensorFlow 2

PyTorch

Activating PyTorch

When a stable Conda package of a framework is released, it's tested and pre-installed on the DLAMI. If you want to run the latest, untested nightly build, you can <u>Install PyTorch's Nightly Build</u> (experimental) manually.

To activate the currently installed framework, follow these instructions on your Deep Learning AMI with Conda.

For PyTorch on Python 3 with CUDA and MKL-DNN, run this command:

```
$ source activate pytorch_p310
```

Start the iPython terminal.

```
(pytorch_p310)$ ipython
```

Run a quick PyTorch program.

```
import torch
x = torch.rand(5, 3)
print(x)
print(x.size())
y = torch.rand(5, 3)
print(torch.add(x, y))
```

You should see the initial random array printed, then its size, and then the addition of another random array.

Install PyTorch's Nightly Build (experimental)

How to install PyTorch from a nightly build

You can install the latest PyTorch build into either or both of the PyTorch Conda environments on your Deep Learning AMI with Conda.

1. • (Option for Python 3) - Activate the Python 3 PyTorch environment:

```
$ source activate pytorch_p310
```
2. The remaining steps assume you are using the pytorch_p310 environment. Remove the currently installed PyTorch:

(pytorch_p310)\$ pip uninstall torch

3. • (Option for GPU instances) - Install the latest nightly build of PyTorch with CUDA.0:

(pytorch_p310)\$ pip install torch_nightly -f https://download.pytorch.org/whl/ nightly/cu100/torch_nightly.html

 (Option for CPU instances) - Install the latest nightly build of PyTorch for instances with no GPUs:

(pytorch_p310)\$ pip install torch_nightly -f https://download.pytorch.org/whl/ nightly/cpu/torch_nightly.html

4. To verify you have successfully installed latest nightly build, start the IPython terminal and check the version of PyTorch.

(pytorch_p310)\$ ipython

```
import torch
print (torch.__version__)
```

The output should print something similar to 1.0.0.dev20180922

5. To verify that the PyTorch nightly build works well with the MNIST example, you can run a test script from PyTorch's examples repository:

```
(pytorch_p310)$ cd ~
(pytorch_p310)$ git clone https://github.com/pytorch/examples.git pytorch_examples
(pytorch_p310)$ cd pytorch_examples/mnist
(pytorch_p310)$ python main.py || exit 1
```

More Tutorials

For further tutorials and examples refer to the framework's official docs, <u>PyTorch documentation</u>, and the <u>PyTorch</u> website.

TensorFlow 2

This tutorial shows how to activate TensorFlow 2 on an instance running the Deep Learning AMI with Conda (DLAMI on Conda) and run a TensorFlow 2 program.

When a stable Conda package of a framework is released, it's tested and pre-installed on the DLAMI.

Activating TensorFlow 2

To run TensorFlow on the DLAMI with Conda

- 1. To activate TensorFlow 2, open an Amazon Elastic Compute Cloud (Amazon EC2) instance of the DLAMI with Conda.
- 2. For TensorFlow 2 and Keras 2 on Python 3 with CUDA 10.1 and MKL-DNN, run this command:

\$ source activate tensorflow2_p310

3. Start the iPython terminal:

(tensorflow2_p310)\$ ipython

4. Run a TensorFlow 2 program to verify that it is working properly:

```
import tensorflow as tf
hello = tf.constant('Hello, TensorFlow!')
tf.print(hello)
```

Hello, TensorFlow! should appear on your screen.

More Tutorials

For more tutorials and examples, see the TensorFlow documentation for the <u>TensorFlow Python</u> <u>API</u> or see the <u>TensorFlow</u> website.

Distributed training using Elastic Fabric Adapter

An <u>Elastic Fabric Adapter</u> (EFA) is a network device that you can attach to your DLAMI instance to accelerate High Performance Computing (HPC) applications. EFA enables you to achieve

the application performance of an on-premises HPC cluster, with the scalability, flexibility, and elasticity provided by the Amazon Cloud.

The following topics show you how to get started using EFA with the DLAMI.

i Note

Choose your DLAMI from this Base GPU DLAMI list

Topics

- Launching a Amazon Deep Learning AMIs Instance With EFA
- Using EFA on the DLAMI

Launching a Amazon Deep Learning AMIs Instance With EFA

The latest Base DLAMI is ready to use with EFA and comes with the required drivers, kernel modules, libfabric, openmpi and the <u>NCCL OFI plugin</u> for GPU instances.

You can find the supported CUDA versions of a Base DLAMI in the release notes.

Note:

• When running a NCCL Application using mpirun on EFA, you will have to specify the full path to the EFA supported installation as:

/opt/amazon/openmpi/bin/mpirun <command>

• To enable your application to use EFA, add FI_PROVIDER="efa" to the mpirun command as shown in Using EFA on the DLAMI.

Topics

- Prepare an EFA Enabled Security Group
- Launch Your Instance
- Verify EFA Attachment

Prepare an EFA Enabled Security Group

EFA requires a security group that allows all inbound and outbound traffic to and from the security group itself. For more information, see the EFA Documentation.

- 1. Open the Amazon EC2 console at <u>https://console.amazonaws.cn/ec2/</u>.
- 2. In the navigation pane, choose Security Groups and then choose Create Security Group.
- 3. In the Create Security Group window, do the following:
 - For **Security group name**, enter a descriptive name for the security group, such as EFAenabled security group.
 - (Optional) For **Description**, enter a brief description of the security group.
 - For **VPC**, select the VPC into which you intend to launch your EFA-enabled instances.
 - Choose Create.
- 4. Select the security group that you created, and on the **Description** tab, copy the **Group ID**.
- 5. On the Inbound and Outbound tabs, do the following:
 - Choose Edit.
 - For Type, choose All traffic.
 - For Source, choose Custom.
 - Paste the security group ID that you copied into the field.
 - Choose Save.
- 6. Enable inbound traffic referring to <u>Authorizing Inbound Traffic for Your Linux Instances</u>. If you skip this step, you won't be able to communicate with your DLAMI instance.

Launch Your Instance

EFA on the Amazon Deep Learning AMIs is currently supported with the following instance types and operating systems:

- P3dn: Amazon Linux 2, Ubuntu 20.04
- P4d, P4de: Amazon Linux 2, Amazon Linux 2023, Ubuntu 20.04, Ubuntu 22.04
- P5, P5e, P5en: Amazon Linux 2, Amazon Linux 2023, Ubuntu 20.04, Ubuntu 22.04

The following section shows how to launch an EFA enabled DLAMI instance. For more information on launching an EFA enabled instance, see <u>Launch EFA-Enabled Instances into a Cluster Placement</u> <u>Group</u>.

- 1. Open the Amazon EC2 console at https://console.amazonaws.cn/ec2/.
- 2. Choose Launch Instance.
- On the Choose an AMI page, select a supported DLAMI found on the <u>DLAMI Release Notes</u> <u>Page</u>
- 4. On the **Choose an Instance Type** page, select one of the following supported instance types and then choose **Next: Configure Instance Details.** Refer to this link for the list of supported instances: Get started with EFA and MPI
- 5. On the **Configure Instance Details** page, do the following:
 - For **Number of instances**, enter the number of EFA-enabled instances that you want to launch.
 - For **Network** and **Subnet**, select the VPC and subnet into which to launch the instances.
 - [Optional] For **Placement group**, select **Add instance to placement group**. For best performance, launch the instances within a placement group.
 - [Optional] For **Placement group name**, select **Add to a new placement group**, enter a descriptive name for the placement group, and then for **Placement group strategy**, select **cluster**.
 - Make sure to enable the **"Elastic Fabric Adapter"** on this page. If this option is disabled, change the subnet to one that supports your selected instance type.
 - In the Network Interfaces section, for device eth0, choose New network interface. You
 can optionally specify a primary IPv4 address and one or more secondary IPv4 addresses. If
 you're launching the instance into a subnet that has an associated IPv6 CIDR block, you can
 optionally specify a primary IPv6 address and one or more secondary IPv6 addresses.
 - Choose Next: Add Storage.
- On the Add Storage page, specify the volumes to attach to the instances in addition to the volumes specified by the AMI (such as the root device volume), and then choose Next: Add Tags.
- 7. On the **Add Tags** page, specify tags for the instances, such as a user-friendly name, and then choose **Next: Configure Security Group**.
- 8. On the **Configure Security Group** page, for **Assign a security group**, select **Select an existing security group**, and then select the security group that you created previously.

9. Choose Review and Launch.

10. On the **Review Instance Launch** page, review the settings, and then choose **Launch** to choose a key pair and to launch your instances.

Verify EFA Attachment

From the Console

After launching the instance, check the instance details in the Amazon Console. To do this, select the instance in the EC2 console and look at the Description Tab in the lower pane on the page. Find the parameter 'Network Interfaces: eth0' and click on eth0 which opens a pop-up. Make sure that 'Elastic Fabric Adapter' is enabled.

If EFA is not enabled, you can fix this by either:

- Terminating the EC2 instance and launching a new one with the same steps. Make sure the EFA is attached.
- Attach EFA to an existing instance.
 - 1. In the EC2 Console, go to Network Interfaces.
 - 2. Click on Create a Network Interface.
 - 3. Select the same subnet that your instance is in.
 - 4. Make sure to enable the 'Elastic Fabric Adapter' and click on Create.
 - 5. Go back to the EC2 Instances Tab and select your instance.
 - 6. Go to Actions: Instance State and stop the instance before you attach EFA.
 - 7. From Actions, select Networking: Attach Network Interface.
 - 8. Select the interface you just created and click on attach.
 - 9. Restart your instance.

From the Instance

The following test script is already present on the DLAMI. Run it to ensure that the kernel modules are loaded correctly.

\$ fi_info -p efa

Your output should look similar to the following.

```
provider: efa
    fabric: EFA-fe80::e5:56ff:fe34:56a8
    domain: efa_0-rdm
    version: 2.0
    type: FI_EP_RDM
    protocol: FI_PROTO_EFA
provider: efa
    fabric: EFA-fe80::e5:56ff:fe34:56a8
    domain: efa_0-dgrm
    version: 2.0
    type: FI_EP_DGRAM
    protocol: FI_PROTO_EFA
provider: efa;ofi_rxd
    fabric: EFA-fe80::e5:56ff:fe34:56a8
    domain: efa_0-dgrm
    version: 1.0
    type: FI_EP_RDM
    protocol: FI_PROTO_RXD
```

Verify Security Group Configuration

The following test script is already present on the DLAMI. Run it to ensure that the security group you created is configured correctly.

```
$ cd /opt/amazon/efa/test/
$ ./efa_test.sh
```

Your output should look similar to the following.

Starting server Starting client											
bytes	#sent	#ack	total	time	MB/sec	usec/xfer	Mxfers/sec				
64	10	=10	1.2k	0.02s	0.06	1123.55	0.00				
256	10	=10	5k	0.00s	17.66	14.50	0.07				
1k	10	=10	20k	0.00s	67.81	15.10	0.07				
4k	10	=10	80k	0.00s	237.45	17.25	0.06				
64k	10	=10	1.2m	0.00s	921.10	71.15	0.01				
1m	10	=10	20m	0.01s	2122.41	494.05	0.00				

If it stops responding or does not complete, ensure that your security group has the correct inbound/outbound rules.

Using EFA on the DLAMI

The following section describes how to use EFA to run multi-node applications on the Amazon Deep Learning AMIs.

Running Multi-Node Applications with EFA

To run an application across a cluster of nodes the following configuration is required

Topics

- Enable Passwordless SSH
- Create Hosts File
- NCCL Tests

Enable Passwordless SSH

Select one node in your cluster as the leader node. The remaining nodes are referred to as the member nodes.

1. On the leader node, generate the RSA keypair.

```
ssh-keygen -t rsa -N "" -f ~/.ssh/id_rsa
```

2. Change the permissions of the private key on the leader node.

chmod 600 ~/.ssh/id_rsa

- Copy the public key ~/.ssh/id_rsa.pub to and append it to ~/.ssh/authorized_keys
 of the member nodes in the cluster.
- 4. You should now be able to directly login to the member nodes from the leader node using the private ip.

```
ssh <member private ip>
```

5. Disable strictHostKeyChecking and enable agent forwarding on the leader node by adding the following to the ~/.ssh/config file on the leader node:

```
Host *
ForwardAgent yes
```

Host * StrictHostKeyChecking no

6. On Amazon Linux 2 instances, run the following command on the leader node to provide correct permissions to the config file:

chmod 600 ~/.ssh/config

Create Hosts File

On the leader node, create a hosts file to identify the nodes in the cluster. The hosts file must have an entry for each node in the cluster. Create a file ~/hosts and add each node using the private ip as follows:

```
localhost slots=8
<private ip of node 1> slots=8
<private ip of node 2> slots=8
```

NCCL Tests

Note

These tests have been run using EFA version 1.38.0 and OFI NCCL Plugin 1.13.2.

Listed below are a subset of NCCL Tests provided by Nvidia to test both functionality and performance over multiple compute nodes

Supported Instances: P3dn, P4, P5, P5e, P5en

Performance Tests

Multi-node NCCL Performance Test on P4d.24xlarge

To check NCCL Performance with EFA, run the standard NCCL Performance test that is available on the official <u>NCCL-Tests Repo</u>. The DLAMI comes with this test already built for CUDA XX.X. You can similarly run your own script with EFA.

When constructing your own script, refer to the following guidance:

- Use the complete path to mpirun as shown in the example while running NCCL applications with EFA.
- Change the params np and N based on the number of instances and GPUs in your cluster.
- Add the NCCL_DEBUG=INFO flag and make sure that the logs indicate EFA usage as "Selected Provider is EFA".
- Set the Training Log Location to parse for validation

```
TRAINING_LOG="testEFA_$(date +"%N").log"
```

Use the command watch nvidia-smi on any of the member nodes to monitor GPU usage. The following watch nvidia-smi commands are for a generic CUDA xx.x version and depend on the Operating System of your instance. You can run the commands for any available CUDA version in your Amazon EC2 instance by replacing the CUDA version in the script.

• Amazon Linux 2, Amazon Linux 2023:

```
$ /opt/amazon/openmpi/bin/mpirun -n 16 -N 8 \
-x NCCL_DEBUG=INFO --mca pml ^cm \
-x LD_LIBRARY_PATH=/usr/local/cuda-xx.x/efa/lib:/usr/local/cuda-xx.x/lib:/usr/
local/cuda-xx.x/lib64:/usr/local/cuda-xx.x:/opt/amazon/efa/lib64:/opt/amazon/openmpi/
lib64:$LD_LIBRARY_PATH \
--hostfile hosts --mca btl tcp,self --mca btl_tcp_if_exclude lo,docker0 --bind-to
none \
/usr/local/cuda-xx.x/efa/test-cuda-xx.x/all_reduce_perf -b 8 -e 16 -f 2 -g 1 -c 1 -n
100 | tee ${TRAINING_LOG}
```

• Ubuntu 20.04, Ubuntu 20.04:

```
$ /opt/amazon/openmpi/bin/mpirun -n 16 -N 8 \
-x NCCL_DEBUG=INFO --mca pml ^cm \
-x LD_LIBRARY_PATH=/usr/local/cuda-xx.x/efa/lib:/usr/local/cuda-xx.x/lib:/usr/
local/cuda-xx.x/lib64:/usr/local/cuda-xx.x:/opt/amazon/efa/lib:/opt/amazon/openmpi/
lib:$LD_LIBRARY_PATH \
--hostfile hosts --mca btl tcp,self --mca btl_tcp_if_exclude lo,docker0 --bind-to
none \
/usr/local/cuda-xx.x/efa/test-cuda-xx.x/all_reduce_perf -b 8 -e 16 -f 2 -g 1 -c 1 -n
100 | tee ${TRAINING_LOG}
```

Your output should look like the following:

```
# nThread 1 nGpus 1 minBytes 8 maxBytes 1073741824 step: 2(factor) warmup iters: 5
iters: 100 agg iters: 1 validation: 1 graph: 0
#
# Using devices
# Rank 0 Group 0 Pid 33378 on ip-172-31-42-25 device 0 [0x10] NVIDIA A100-
SXM4-40GB
# Rank 1 Group 0 Pid 33379 on ip-172-31-42-25 device 1 [0x10] NVIDIA A100-
SXM4-40GB
# Rank 2 Group 0 Pid 33380 on ip-172-31-42-25 device 2 [0x20] NVIDIA A100-
SXM4-40GB
# Rank 3 Group 0 Pid 33381 on ip-172-31-42-25 device 3 [0x20] NVIDIA A100-
SXM4-40GB
# Rank 4 Group 0 Pid 33382 on ip-172-31-42-25 device 4 [0x90] NVIDIA A100-
SXM4-40GB
# Rank 5 Group 0 Pid 33383 on ip-172-31-42-25 device 5 [0x90] NVIDIA A100-
SXM4-40GB
# Rank 6 Group 0 Pid 33384 on ip-172-31-42-25 device 6 [0xa0] NVIDIA A100-
SXM4-40GB
#
  Rank 7 Group 0 Pid 33385 on ip-172-31-42-25 device 7 [0xa0] NVIDIA A100-
SXM4-40GB
#
  Rank 8 Group 0 Pid 30378 on ip-172-31-43-8 device 0 [0x10] NVIDIA A100-SXM4-40GB
#
  Rank 9 Group 0 Pid
                        30379 on ip-172-31-43-8 device 1 [0x10] NVIDIA A100-SXM4-40GB
#
  Rank 10 Group 0 Pid 30380 on ip-172-31-43-8 device 2 [0x20] NVIDIA A100-SXM4-40GB
  Rank 11 Group 0 Pid 30381 on ip-172-31-43-8 device 3 [0x20] NVIDIA A100-SXM4-40GB
#
#
  Rank 12 Group 0 Pid 30382 on ip-172-31-43-8 device 4 [0x90] NVIDIA A100-SXM4-40GB
#
  Rank 13 Group 0 Pid 30383 on ip-172-31-43-8 device 5 [0x90] NVIDIA A100-SXM4-40GB
  Rank 14 Group 0 Pid 30384 on ip-172-31-43-8 device 6 [0xa0] NVIDIA A100-SXM4-40GB
#
 Rank 15 Group 0 Pid 30385 on ip-172-31-43-8 device 7 [0xa0] NVIDIA A100-SXM4-40GB
#
ip-172-31-42-25:33385:33385 [7] NCCL INFO cudaDriverVersion 12060
ip-172-31-43-8:30383:30383 [5] NCCL INFO Bootstrap : Using ens32:172.31.43.8
ip-172-31-43-8:30383:30383 [5] NCCL INFO NCCL version 2.23.4+cuda12.5
. . .
ip-172-31-42-25:33384:33451 [6] NCCL INFO NET/OFI Initializing aws-ofi-nccl 1.13.2-aws
ip-172-31-42-25:33384:33451 [6] NCCL INFO NET/OFI Using Libfabric version 1.22
ip-172-31-42-25:33384:33451 [6] NCCL INFO NET/OFI Using CUDA driver version 12060 with
runtime 12050
ip-172-31-42-25:33384:33451 [6] NCCL INFO NET/OFI Configuring AWS-specific options
ip-172-31-42-25:33384:33451 [6] NCCL INFO NET/OFI Setting provider_filter to efa
ip-172-31-42-25:33384:33451 [6] NCCL INFO NET/OFI Setting FI_EFA_FORK_SAFE environment
variable to 1
ip-172-31-42-25:33384:33451 [6] NCCL INFO NET/OFI Setting NCCL_NVLSTREE_MAX_CHUNKSIZE
to 512KiB
```

Elastic Fabric Adapter

••				~		out trunc	atad				
				S	ome out	out trunc	ated		out-of-	place	
	in-p	lace									
<u>:</u>	size	COL	unt		type	redop	root	time	algbw	busbw	#wrong
time	algbw busbw #wrong										
÷	(B)	(element	ts)					(us)	(GB/s)	(GB/s)	
(us)	(GB/s)	(GB/s)									
	8		2		float	sum	-1	180.3	0.00	0.00	(
179.3	0.00	0.00		0							
	16		4		float	sum	-1	178.1	0.00	0.00	(
177.6	0.00	0.00	-	0							
477 0	32		8	•	float	sum	-1	178.5	0.00	0.00	(
177.9	0.00	0.00	10	0	61		1	170 0	0 00	0 00	
170 7	64	0 00	16	0	float	sum	-1	178.8	0.00	0.00	(
178.7	0.00 128	0.00	32	Ø	float	cum	-1	178.2	0.00	0.00	(
177.8	0.00	0.00	52	0	Tituat	sum	-1	170.2	0.00	0.00	,
1//.0	256	0.00	64	U	float	sum	-1	178.6	0.00	0.00	(
178.8	0.00	0.00	04	0	Tiout	Sum	-	1/0.0	0.00	0.00	
2/010	512		128	Ū	float	sum	-1	177.2	0.00	0.01	(
177.1	0.00	0.01		0							
	1024		256		float	sum	-1	179.2	0.01	0.01	(
179.3	0.01	0.01		0							
	2048	Į	512		float	sum	-1	181.3	0.01	0.02	(
181.2	0.01	0.02		0							
	4096	10	024		float	sum	-1	184.2	0.02	0.04	(
183.9	0.02	0.04		0							
	8192	20	048		float	sum	-1	191.2	0.04	0.08	(
190.6	0.04	0.08		0							
	16384					sum	-1	202.5	0.08	0.15	(
		0.15									
						sum	-1	233.0	0.14	0.26	
		0.26									
	65536				float	sum	-1	238.6	0.27	0.51	(
	0.28				£ 1 +		1	277 2		1 0/	
						sum	-1	237.2	0.55	1.04	(
		1.04					1		1 00	1 00	
	1.06					sum	-1	248.3	1.00	1.98	(

ip-172-31-42-25:33384:33451 [6] NCCL INFO NET/OFI Setting NCCL_NVLS_CHUNKSIZE to 512KiB ip-172-31-42-25:33384:33451 [6] NCCL INFO NET/OFI Running on p4d.24xlarge platform, Amazon Deep Learning AMIs

52	4288	131072		float	sum	-1	309.2	1.70	3.18	0
307.7	1.70	3.20	0							
104	8576	262144		float	sum	-1	408.7	2.57	4.81	0
404.3	2.59	4.86	0							
209	7152	524288		float	sum	-1	613.5	3.42	6.41	0
607.9	3.45	6.47	0							
419	4304	1048576		float	sum	-1	924.5	4.54	8.51	0
914.8	4.58	8.60	0							
838	8608	2097152		float	sum	-1	1059.5	7.92	14.85	0
1054.3	7.96	14.92	0							
1677	7216	4194304		float	sum	-1	1269.9	13.21	24.77	0
1272.0	13.19	24.73	0							
3355	4432	8388608		float	sum	-1	1642.7	20.43	38.30	0
1636.7	20.50	38.44	0							
6710	8864	16777216		float	sum	-1	2446.7	27.43	51.43	0
	27.44									
13421	7728	33554432		float	sum	-1	4143.6	32.39	60.73	0
		60.75								
		67108864			sum	-1	7351.9	36.51	68.46	0
		68.51								
		134217728			sum	-1	13717	39.14	73.39	0
		73.46								
107374	1824	268435456		float	sum	-1	26416	40.65	76.21	0
26420	40.64	76.20	0							
•••										
# Out of bounds values : 0 OK										
# Avg bus bandwidth : 15.5514										

Validation Tests

To Validate that the EFA tests returned a valid result, please use the following tests to confirm:

• Get the instance type using EC2 Instance Metadata:

```
TOKEN=$(curl -X PUT "http://169.254.169.254/latest/api/token" -H "X-aws-ec2-metadata-
token-ttl-seconds: 21600")
INSTANCE_TYPE=$(curl -H "X-aws-ec2-metadata-token: $TOKEN" -v http://169.254.169.254/
latest/meta-data/instance-type)
```

- Run the Performance Tests
- Set the Following Parameters

CUDA_VERSION

CUDA_RUNTIME_VERSION NCCL_VERSION

• Validate the Results as shown:

```
RETURN_VAL=`echo $?`
if [ ${RETURN_VAL} -eq 0 ]; then
    # [0] NCCL INFO NET/OFI Initializing aws-ofi-nccl 1.13.2-aws
    # [0] NCCL INFO NET/OFI Using CUDA driver version 12060 with runtime 12010
    # cudaDriverVersion 12060 --> This is max supported cuda version by nvidia
 driver
    # NCCL version 2.23.4+cuda12.5 --> This is NCCL version compiled with cuda
 version
    # Validation of logs
    grep "NET/OFI Configuring AWS-specific options" ${TRAINING_LOG} || { echo "AWS-
specific options text not found"; exit 1; }
    grep "busbw" ${TRAINING_LOG} || { echo "busbw text not found"; exit 1; }
    grep "Avg bus bandwidth " ${TRAINING_LOG} || { echo "Avg bus bandwidth text not
 found"; exit 1; }
    grep "NCCL version $NCCL_VERSION" ${TRAINING_LOG} || { echo "Text not found: NCCL
 version $NCCL_VERSION"; exit 1; }
    if [[ ${INSTANCE_TYPE} == "p4d.24xlarge" ]]; then
        grep "NET/Libfabric/0/GDRDMA" ${TRAINING_LOG} || { echo "Text not found: NET/
Libfabric/0/GDRDMA"; exit 1; }
        grep "NET/OFI Selected Provider is efa (found 4 nics)" ${TRAINING_LOG} ||
 { echo "Selected Provider is efa text not found"; exit 1; }
    elif [[ ${INSTANCE_TYPE} == "p4de.24xlarge" ]]; then
        grep "NET/Libfabric/0/GDRDMA" ${TRAINING_LOG} || { echo "Avg bus bandwidth
 text not found"; exit 1; }
        grep "NET/OFI Selected Provider is efa (found 4 nics)" ${TRAINING_LOG} ||
 { echo "Avg bus bandwidth text not found"; exit 1; }
    elif [[ ${INSTANCE_TYPE} == "p5.48xlarge" ]]; then
        grep "NET/Libfabric/0/GDRDMA" ${TRAINING_LOG} || { echo "Avg bus bandwidth
 text not found"; exit 1; }
        grep "NET/OFI Selected Provider is efa (found 32 nics)" ${TRAINING_LOG} ||
 { echo "Avg bus bandwidth text not found"; exit 1; }
    elif [[ ${INSTANCE_TYPE} == "p5e.48xlarge" ]]; then
        grep "NET/Libfabric/0/GDRDMA" ${TRAINING_LOG} || { echo "Avg bus bandwidth
 text not found"; exit 1; }
        grep "NET/OFI Selected Provider is efa (found 32 nics)" ${TRAINING_LOG} ||
 { echo "Avg bus bandwidth text not found"; exit 1; }
```

 To access the benchmark data, we can parse the final row of table output from the Multi Node all_reduce test:

```
benchmark=$(sudo cat ${TRAINING_LOG} | grep '1073741824' | tail -n1 | awk -F " "
  '{{print $12}}' | sed 's/ //' | sed 's/ 5e-07//')
if [[ -z "${benchmark}" ]]; then
  echo "benchmark variable is empty"
  exit 1
fi
echo "Benchmark throughput: ${benchmark}"
```

GPU Monitoring and Optimization

The following section will guide you through GPU optimization and monitoring options. This section is organized like a typical workflow with monitoring overseeing preprocessing and training.

- Monitoring
 - Monitor GPUs with CloudWatch
- Optimization
 - Preprocessing
 - Training

Monitoring

Your DLAMI comes preinstalled with several GPU monitoring tools. This guide also mentions tools that are available to download and install.

- <u>Monitor GPUs with CloudWatch</u> a preinstalled utility that reports GPU usage statistics to Amazon CloudWatch.
- <u>nvidia-smi CLI</u> a utility to monitor overall GPU compute and memory utilization. This is preinstalled on your Amazon Deep Learning AMIs (DLAMI).
- <u>NVML C library</u> a C-based API to directly access GPU monitoring and management functions. This used by the nvidia-smi CLI under the hood and is preinstalled on your DLAMI. It also has Python and Perl bindings to facilitate development in those languages. The gpumon.py utility preinstalled on your DLAMI uses the pynvml package from <u>nvidia-ml-py</u>.
- <u>NVIDIA DCGM</u> A cluster management tool. Visit the developer page to learn how to install and configure this tool.

🚺 Tip

Check out NVIDIA's developer blog for the latest info on using the CUDA tools installed your DLAMI:

• Monitoring TensorCore utilization using Nsight IDE and nvprof.

Monitor GPUs with CloudWatch

When you use your DLAMI with a GPU you might find that you are looking for ways to track its usage during training or inference. This can be useful for optimizing your data pipeline, and tuning your deep learning network.

There are two ways to configure GPU metrics with CloudWatch:

- Configure metrics with the Amazon CloudWatch agent (Recommended)
- Configure metrics with the preinstalled gpumon.py script

Configure metrics with the Amazon CloudWatch agent (Recommended)

Integrate your DLAMI with the <u>unified CloudWatch agent</u> to configure GPU metrics and monitor the utilization of GPU coprocesses in Amazon EC2 accelerated instances.

There are four ways to configure <u>GPU metrics</u> with your DLAMI:

- <u>Configure minimal GPU metrics</u>
- <u>Configure partial GPU metrics</u>
- <u>Configure all available GPU metrics</u>
- <u>Configure custom GPU metrics</u>

For information on updates and security patches, see <u>Security patching for the Amazon</u> <u>CloudWatch agent</u>

Prerequisites

To get started, you must configure Amazon EC2 instance IAM permissions that allow your instance to push metrics to CloudWatch. For detailed steps, see <u>Create IAM roles and users for use with the</u> <u>CloudWatch agent</u>.

Configure minimal GPU metrics

Configure minimal GPU metrics using the dlami-cloudwatch-agent@minimal systemd service. This service configures the following metrics:

- utilization_gpu
- utilization_memory

You can find the systemd service for minimal preconfigured GPU metrics in the following location:

/opt/aws/amazon-cloudwatch-agent/etc/dlami-amazon-cloudwatch-agent-minimal.json

Enable and start the systemd service with the following commands:

```
sudo systemctl enable dlami-cloudwatch-agent@minimal
sudo systemctl start dlami-cloudwatch-agent@minimal
```

Configure partial GPU metrics

Configure partial GPU metrics using the dlami-cloudwatch-agent@partial systemd service. This service configures the following metrics:

- utilization_gpu
- utilization_memory
- memory_total
- memory_used
- memory_free

You can find the systemd service for partial preconfigured GPU metrics in the following location:

/opt/aws/amazon-cloudwatch-agent/etc/dlami-amazon-cloudwatch-agent-partial.json

Enable and start the systemd service with the following commands:

```
sudo systemctl enable dlami-cloudwatch-agent@partial
sudo systemctl start dlami-cloudwatch-agent@partial
```

Configure all available GPU metrics

Configure all available GPU metrics using the dlami-cloudwatch-agent@all systemd service. This service configures the following metrics:

- utilization_gpu
- utilization_memory
- memory_total
- memory_used
- memory_free
- temperature_gpu
- power_draw
- fan_speed
- pcie_link_gen_current
- pcie_link_width_current

- encoder_stats_session_count
- encoder_stats_average_fps
- encoder_stats_average_latency
- clocks_current_graphics
- clocks_current_sm
- clocks_current_memory
- clocks_current_video

You can find the systemd service for all available preconfigured GPU metrics in the following location:

/opt/aws/amazon-cloudwatch-agent/etc/dlami-amazon-cloudwatch-agent-all.json

Enable and start the systemd service with the following commands:

```
sudo systemctl enable dlami-cloudwatch-agent@all
sudo systemctl start dlami-cloudwatch-agent@all
```

Configure custom GPU metrics

If the preconfigured metrics do not meet your requirements, you can create a custom CloudWatch agent configuration file.

Create a custom configuration file

To create a custom configuration file, refer to the detailed steps in <u>Manually create or edit the</u> CloudWatch agent configuration file.

For this example, assume that the schema definition is located at /opt/aws/amazoncloudwatch-agent/etc/amazon-cloudwatch-agent.json.

Configure metrics with your custom file

Run the following command to configure the CloudWatch agent according to your custom file:

```
sudo /opt/aws/amazon-cloudwatch-agent/bin/amazon-cloudwatch-agent-ctl \
-a fetch-config -m ec2 -s -c \
file:/opt/aws/amazon-cloudwatch-agent/etc/amazon-cloudwatch-agent.json
```

Security patching for the Amazon CloudWatch agent

Newly released DLAMIs are configured with the latest available Amazon CloudWatch agent security patches. Refer to the following sections to update your current DLAMI with the latest security patches depending on your operating system of choice.

Amazon Linux 2

Use yum to get the latest Amazon CloudWatch agent security patches for an Amazon Linux 2 DLAMI.

sudo yum update

Ubuntu

To get the latest Amazon CloudWatch security patches for a DLAMI with Ubuntu, it is necessary to reinstall the Amazon CloudWatch agent using an Amazon S3 download link.

wget https://s3.region.amazonaws.com/amazoncloudwatch-agent-region/ubuntu/arm64/latest/ amazon-cloudwatch-agent.deb

For more information on installing the Amazon CloudWatch agent using Amazon S3 download links, see Installing and running the CloudWatch agent on your servers.

Configure metrics with the preinstalled gpumon.py script

A utility called gpumon.py is preinstalled on your DLAMI. It integrates with CloudWatch and supports monitoring of per-GPU usage: GPU memory, GPU temperature, and GPU Power. The script periodically sends the monitored data to CloudWatch. You can configure the level of granularity for data being sent to CloudWatch by changing a few settings in the script. Before starting the script, however, you will need to setup CloudWatch to receive the metrics.

How to setup and run GPU monitoring with CloudWatch

 Create an IAM user, or modify an existing one to have a policy for publishing the metric to CloudWatch. If you create a new user please take note of the credentials as you will need these in the next step.

The IAM policy to search for is "cloudwatch:PutMetricData". The policy that is added is as follows:

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Action": [
                "cloudwatch:PutMetricData"
        ],
            "Effect": "Allow",
            "Resource": "*"
        }
    ]
}
```

🚺 Tip

For more information on creating an IAM user and adding policies for CloudWatch, refer to the <u>CloudWatch documentation</u>.

2. On your DLAMI, run Amazon configure and specify the IAM user credentials.

\$ aws configure

3. You might need to make some modifications to the gpumon utility before you run it. You can find the gpumon utility and README in the location defined in the following code block. For more information on the gpumon.py script, see the Amazon S3 location of the script.

Options:

- Change the region in gpumon.py if your instance is NOT in us-east-1.
- Change other parameters such as the CloudWatch namespace or the reporting period with store_reso.
- 4. Currently the script only supports Python 3. Activate your preferred framework's Python 3 environment or activate the DLAMI general Python 3 environment.

\$ source activate python3

5. Run the gpumon utility in background.

(python3)\$ python gpumon.py &

 Open your browser to the <u>https://console.amazonaws.cn/cloudwatch/</u> then select metric. It will have a namespace 'DeepLearningTrain'.

(i) Tip You can change the namespace by modifying gpumon.py. You can also modify the reporting interval by adjusting store_reso.

The following is an example CloudWatch chart reporting on a run of gpumon.py monitoring a training job on p2.8xlarge instance.



You might be interested in these other topics on GPU monitoring and optimization:

- Monitoring
 - Monitor GPUs with CloudWatch
- Optimization

- Preprocessing
- Training

Optimization

To make the most of your GPUs, you can optimize your data pipeline and tune your deep learning network. As the following chart describes, a naive or basic implementation of a neural network might use the GPU inconsistently and not to its fullest potential. When you optimize your preprocessing and data loading, you can reduce the bottleneck from your CPU to your GPU. You can adjust the neural network itself, by using hybridization (when supported by the framework), adjusting batch size, and synchronizing calls. You can also use multiple-precision (float16 or int8) training in most frameworks, which can have a dramatic effect on improving throughput.

The following chart shows the cumulative performance gains when applying different optimizations. Your results will depend on the data you are processing and the network you are optimizing.



Example GPU performance optimizations. Chart source: Performance Tricks with MXNet Gluon

The following guides introduce options that will work with your DLAMI and help you boost GPU performance.

Topics

- Preprocessing
- Training

Preprocessing

Data preprocessing through transformations or augmentations can often be a CPU-bound process, and this can be the bottleneck in your overall pipeline. Frameworks have built-in operators for image processing, but DALI (Data Augmentation Library) demonstrates improved performance over frameworks' built-in options.

- NVIDIA Data Augmentation Library (DALI): DALI offloads data augmentation to the GPU. It
 is not preinstalled on the DLAMI, but you can access it by installing it or loading a supported
 framework container on your DLAMI or other Amazon Elastic Compute Cloud instance. Refer
 to the <u>DALI project page</u> on the NVIDIA website for details. For an example use-case and to
 download code samples, see the <u>SageMaker Preprocessing Training Performance</u> sample.
- nvJPEG: a GPU-accelerated JPEG decoder library for C programmers. It supports decoding single images or batches as well as subsequent transformation operations that are common in deep learning. nvJPEG comes built-in with DALI, or you can download from the <u>NVIDIA website's</u> <u>nvjpeg page</u> and use it separately.

You might be interested in these other topics on GPU monitoring and optimization:

- Monitoring
 - Monitor GPUs with CloudWatch
- Optimization
 - Preprocessing
 - Training

Training

With mixed-precision training you can deploy larger networks with the same amount of memory, or reduce memory usage compared to your single or double precision network, and you will see compute performance increases. You also get the benefit of smaller and faster data transfers, an important factor in multiple node distributed training. To take advantage of mixed-precision training you need to adjust data casting and loss scaling. The following are guides describing how to do this for the frameworks that support mixed-precision.

• <u>NVIDIA Deep Learning SDK</u> - docs on the NVIDIA website describing mixed-precision implementation for MXNet, PyTorch, and TensorFlow.

🚺 Tip

Be sure to check the website for your framework of choice, and search for "mixed precision" or "fp16" for the latest optimization techniques. Here are some mixed-precision guides you might find helpful:

- Mixed-precision training with TensorFlow (video) on the NVIDIA blog site.
- <u>Mixed-precision training using float16 with MXNet</u> an FAQ article on the MXNet website.
- <u>NVIDIA Apex: a tool for easy mixed-precision training with PyTorch</u> a blog article on the NVIDIA website.

You might be interested in these other topics on GPU monitoring and optimization:

- Monitoring
 - Monitor GPUs with CloudWatch
- Optimization
 - Preprocessing
 - Training

The Amazon Inferentia Chip With DLAMI

Amazon Inferentia is a custom machine learning chip designed by Amazon that you can use for high-performance inference predictions. In order to use the chip, set up an Amazon Elastic Compute Cloud instance and use the Amazon Neuron software development kit (SDK) to invoke the Inferentia chip. To provide customers with the best Inferentia experience, Neuron has been built into the Amazon Deep Learning AMIs (DLAMI).

The following topics show you how to get started using Inferentia with the DLAMI.

Contents

- Launching a DLAMI Instance with Amazon Neuron
- Using the DLAMI with Amazon Neuron

Launching a DLAMI Instance with Amazon Neuron

The latest DLAMI is ready to use with Amazon Inferentia and comes with the Amazon Neuron API package. To launch a DLAMI instance, see <u>Launching and Configuring a DLAMI</u>. After you have a DLAMI, use the steps here to ensure that your Amazon Inferentia chip and Amazon Neuron resources are active.

Contents

- Verify Your Instance
- Identifying Amazon Inferentia Devices
- View Resource Usage
- Using Neuron Monitor (neuron-monitor)
- Upgrading Neuron Software

Verify Your Instance

Before using your instance, verify that it's properly setup and configured with Neuron.

Identifying Amazon Inferentia Devices

To identify the number of Inferentia devices on your instance, use the following command:

```
neuron-ls
```

If your instance has Inferentia devices attached to it, your output will look similar to the following:

| NEURON | NEURON | NEURON | CONNECTED | PCI I | DEVICE | CORES | MEMORY | DEVICES | BDF 0 | 4 8 GB | 1 | 0000:00:1c.0 | | 1 | 4 | 8 GB 2,0 | 0000:00:1d.0 | | 2 | 4 | 8 GB | 3, 1 | 0000:00:1e.0 | | 2 | 0000:00:1f.0 | | 3 | 4 8 GB

The supplied output is taken from an Inf1.6xlarge instance and includes the following columns:

- NEURON DEVICE: The logical ID assigned to the NeuronDevice. This ID is used when configuring multiple runtimes to use different NeuronDevices.
- NEURON CORES: The number of NeuronCores present in the NeuronDevice.
- NEURON MEMORY: The amount of DRAM memory in the NeuronDevice.
- CONNECTED DEVICES: Other NeuronDevices connected to the NeuronDevice.
- PCI BDF: The PCI Bus Device Function (BDF) ID of the NeuronDevice.

View Resource Usage

View useful information about NeuronCore and vCPU utilization, memory usage, loaded models, and Neuron applications with the neuron-top command. Launching neuron-top with no arguments will show data for all machine learning applications that utilize NeuronCores.

neuron-top

When an application is using four NeuronCores, the output should look similar to the following image:

neuron-top									
ND3	NC1 [[109%] [0.08%] [0.08%] [0.08%]			NC3 					
<u>VCPU and Memory Info</u> System vCPU Usage Runtime Memory Host		Runtime vCPU Usage Runtime Memory Device 198.3MB							
Loaded Models									
	м	odel ID	Host Memory 2.5MB	Device Memory 198.3MB					
integ-tests/out-test7_resnet50_v2_fp16_b1_tpb1_tf [+] NG1 [+] NG2 [+] NG3	1	0001	638.5KB 638.5KB 638.5KB 638.5KB	49.6MB 49.8MB 49.8MB 49.8MB					
Neuron Apps >{1}tinference and is q: quit arrows: move tree sele	<pre>[2]:inference app 2 ction enter: expand/collapse tree item</pre>	<pre>[3]:inference app 3 x: expand/collapse entire tree</pre>	<pre>[4]:inference app 4 a/d: previous/next tab</pre>	1-9: select tab					

For more information on resources to monitor and optimize Neuron-based inference applications, see Neuron Tools.

Using Neuron Monitor (neuron-monitor)

Neuron Monitor collects metrics from the Neuron runtimes running on the system and streams the collected data to stdout in JSON format. These metrics are organized into metric groups that you configure by providing a configuration file. For more information on Neuron Monitor, see the <u>User</u> Guide for Neuron Monitor.

Upgrading Neuron Software

For information on how to update Neuron SDK software within DLAMI, see the Amazon Neuron Setup Guide.

Next Step

Using the DLAMI with Amazon Neuron

Using the DLAMI with Amazon Neuron

A typical workflow with the Amazon Neuron SDK is to compile a previously trained machine learning model on a compilation server. After this, distribute the artifacts to the Inf1 instances for execution. Amazon Deep Learning AMIs (DLAMI) comes pre-installed with everything you need to compile and run inference in an Inf1 instance that uses Inferentia.

The following sections describe how to use the DLAMI with Inferentia.

Contents

- Using TensorFlow-Neuron and the Amazon Neuron Compiler
- Using Amazon Neuron TensorFlow Serving
- Using MXNet-Neuron and the Amazon Neuron Compiler
- Using MXNet-Neuron Model Serving
- Using PyTorch-Neuron and the Amazon Neuron Compiler

Using TensorFlow-Neuron and the Amazon Neuron Compiler

This tutorial shows how to use the Amazon Neuron compiler to compile the Keras ResNet-50 model and export it as a saved model in SavedModel format. This format is a typical TensorFlow model

interchangeable format. You also learn how to run inference on an Inf1 instance with example input.

For more information about the Neuron SDK, see the Amazon Neuron SDK documentation.

Contents

- Prerequisites
- Activate the Conda environment
- <u>Resnet50 Compilation</u>
- <u>ResNet50 Inference</u>

Prerequisites

Before using this tutorial, you should have completed the set up steps in <u>Launching a DLAMI</u> <u>Instance with Amazon Neuron</u>. You should also have a familiarity with deep learning and using the DLAMI.

Activate the Conda environment

Activate the TensorFlow-Neuron conda environment using the following command:

```
source activate aws_neuron_tensorflow_p36
```

To exit the current conda environment, run the following command:

source deactivate

Resnet50 Compilation

Create a Python script called **tensorflow_compile_resnet50.py** that has the following content. This Python script compiles the Keras ResNet50 model and exports it as a saved model.

```
import os
import time
import shutil
import tensorflow as tf
```

```
import tensorflow.neuron as tfn
import tensorflow.compat.v1.keras as keras
from tensorflow.keras.applications.resnet50 import ResNet50
from tensorflow.keras.applications.resnet50 import preprocess_input
# Create a workspace
WORKSPACE = './ws_resnet50'
os.makedirs(WORKSPACE, exist_ok=True)
# Prepare export directory (old one removed)
model_dir = os.path.join(WORKSPACE, 'resnet50')
compiled_model_dir = os.path.join(WORKSPACE, 'resnet50_neuron')
shutil.rmtree(model_dir, ignore_errors=True)
shutil.rmtree(compiled_model_dir, ignore_errors=True)
# Instantiate Keras ResNet50 model
keras.backend.set_learning_phase(0)
model = ResNet50(weights='imagenet')
# Export SavedModel
tf.saved_model.simple_save(
 session
                    = keras.backend.get_session(),
 export_dir
                    = model_dir,
                    = {'input': model.inputs[0]},
 inputs
                    = {'output': model.outputs[0]})
 outputs
# Compile using Neuron
tfn.saved_model.compile(model_dir, compiled_model_dir)
# Prepare SavedModel for uploading to Inf1 instance
shutil.make_archive(compiled_model_dir, 'zip', WORKSPACE, 'resnet50_neuron')
```

Compile the model using the following command:

python tensorflow_compile_resnet50.py

The compilation process will take a few minutes. When it completes, your output should look like the following:

INFO:tensorflow:fusing subgraph neuron_op_d6f098c01c780733 with neuron-cc

```
INF0:tensorflow:Number of operations in TensorFlow session: 4638
INF0:tensorflow:Number of operations after tf.neuron optimizations: 556
INF0:tensorflow:Number of operations placed on Neuron runtime: 554
INF0:tensorflow:Successfully converted ./ws_resnet50/resnet50 to ./ws_resnet50/
resnet50_neuron
```

•••

After compilation, the saved model is zipped at **ws_resnet50/resnet50_neuron.zip**. Unzip the model and download the sample image for inference using the following commands:

```
unzip ws_resnet50/resnet50_neuron.zip -d .
curl -0 https://raw.githubusercontent.com/awslabs/mxnet-model-server/master/docs/
images/kitten_small.jpg
```

ResNet50 Inference

Create a Python script called **tensorflow_infer_resnet50.py** that has the following content. This script runs inference on the downloaded model using a previously compiled inference model.

```
import os
import numpy as np
import tensorflow as tf
from tensorflow.keras.preprocessing import image
from tensorflow.keras.applications import resnet50
# Create input from image
img_sgl = image.load_img('kitten_small.jpg', target_size=(224, 224))
img_arr = image.img_to_array(img_sql)
img_arr2 = np.expand_dims(img_arr, axis=0)
img_arr3 = resnet50.preprocess_input(img_arr2)
# Load model
COMPILED_MODEL_DIR = './ws_resnet50/resnet50_neuron/'
predictor_inferentia = tf.contrib.predictor.from_saved_model(COMPILED_MODEL_DIR)
# Run inference
model_feed_dict={'input': img_arr3}
infa_rslts = predictor_inferentia(model_feed_dict);
# Display results
print(resnet50.decode_predictions(infa_rslts["output"], top=5)[0])
```

Run inference on the model using the following command:

python tensorflow_infer_resnet50.py

Your output should look like the following:

[('n02123045', 'tabby', 0.6918919), ('n02127052', 'lynx', 0.12770271), ('n02123159', 'tiger_cat', 0.08277027), ('n02124075', 'Egyptian_cat', 0.06418919), ('n02128757', 'snow_leopard', 0.009290541)]

Next Step

Using Amazon Neuron TensorFlow Serving

Using Amazon Neuron TensorFlow Serving

This tutorial shows how to construct a graph and add an Amazon Neuron compilation step before exporting the saved model to use with TensorFlow Serving. TensorFlow Serving is a serving system that allows you to scale-up inference across a network. Neuron TensorFlow Serving uses the same API as normal TensorFlow Serving. The only difference is that a saved model must be compiled for Amazon Inferentia and the entry point is a different binary named tensorflow_model_server_neuron. The binary is found at /usr/local/bin/ tensorflow_model_server_neuron and is pre-installed in the DLAMI.

For more information about the Neuron SDK, see the Amazon Neuron SDK documentation.

Contents

- Prerequisites
- Activate the Conda environment
- Compile and Export the Saved Model
- Serving the Saved Model
- Generate inference requests to the model server

Prerequisites

Before using this tutorial, you should have completed the set up steps in <u>Launching a DLAMI</u> <u>Instance with Amazon Neuron</u>. You should also have a familiarity with deep learning and using the DLAMI.

Activate the Conda environment

Activate the TensorFlow-Neuron conda environment using the following command:

source activate aws_neuron_tensorflow_p36

If you need to exit the current conda environment, run:

source deactivate

Compile and Export the Saved Model

Create a Python script called tensorflow-model-server-compile.py with the following content. This script constructs a graph and compiles it using Neuron. It then exports the compiled graph as a saved model.

```
import tensorflow as tf
import tensorflow.neuron
import os
tf.keras.backend.set_learning_phase(0)
model = tf.keras.applications.ResNet50(weights='imagenet')
sess = tf.keras.backend.get_session()
inputs = {'input': model.inputs[0]}
outputs = {'output': model.outputs[0]}
# save the model using tf.saved_model.simple_save
modeldir = "./resnet50/1"
tf.saved_model.simple_save(sess, modeldir, inputs, outputs)
# compile the model for Inferentia
neuron_modeldir = os.path.join(os.path.expanduser('~'), 'resnet50_inf1', '1')
tf.neuron.saved_model.compile(modeldir, neuron_modeldir, batch_size=1)
```

Compile the model using the following command:

python tensorflow-model-server-compile.py

Your output should look like the following:

```
...
INFO:tensorflow:fusing subgraph neuron_op_d6f098c01c780733 with neuron-cc
INFO:tensorflow:Number of operations in TensorFlow session: 4638
INFO:tensorflow:Number of operations after tf.neuron optimizations: 556
INFO:tensorflow:Number of operations placed on Neuron runtime: 554
INFO:tensorflow:Successfully converted ./resnet50/1 to /home/ubuntu/resnet50_inf1/1
```

Serving the Saved Model

Once the model has been compiled, you can use the following command to serve the saved model with the tensorflow_model_server_neuron binary:

```
tensorflow_model_server_neuron --model_name=resnet50_inf1 \
          --model_base_path=$HOME/resnet50_inf1/ --port=8500 &
```

Your output should look like the following. The compiled model is staged in the Inferentia device's DRAM by the server to prepare for inference.

```
...
2019-11-22 01:20:32.075856: I external/org_tensorflow/tensorflow/cc/saved_model/
loader.cc:311] SavedModel load for tags { serve }; Status: success. Took 40764
microseconds.
2019-11-22 01:20:32.075888: I tensorflow_serving/servables/tensorflow/
saved_model_warmup.cc:105] No warmup data file found at /home/ubuntu/resnet50_inf1/1/
assets.extra/tf_serving_warmup_requests
2019-11-22 01:20:32.075950: I tensorflow_serving/core/loader_harness.cc:87]
Successfully loaded servable version {name: resnet50_inf1 version: 1}
2019-11-22 01:20:32.077859: I tensorflow_serving/model_servers/
server.cc:353] Running gRPC ModelServer at 0.0.0.0:8500 ...
```

Generate inference requests to the model server

Create a Python script called tensorflow-model-server-infer.py with the following content. This script runs inference via gRPC, which is service framework.

```
import numpy as np
```

```
import grpc
import tensorflow as tf
from tensorflow.keras.preprocessing import image
from tensorflow.keras.applications.resnet50 import preprocess_input
from tensorflow_serving.apis import predict_pb2
from tensorflow_serving.apis import prediction_service_pb2_grpc
from tensorflow.keras.applications.resnet50 import decode_predictions
if __name__ == '__main__':
    channel = grpc.insecure_channel('localhost:8500')
    stub = prediction_service_pb2_grpc.PredictionServiceStub(channel)
    img_file = tf.keras.utils.get_file(
        "./kitten_small.jpg",
        "https://raw.githubusercontent.com/awslabs/mxnet-model-server/master/docs/
images/kitten_small.jpg")
    img = image.load_img(img_file, target_size=(224, 224))
    img_array = preprocess_input(image.img_to_array(img)[None, ...])
    request = predict_pb2.PredictRequest()
    request.model_spec.name = 'resnet50_inf1'
    request.inputs['input'].CopyFrom(
        tf.contrib.util.make_tensor_proto(img_array, shape=img_array.shape))
    result = stub.Predict(request)
    prediction = tf.make_ndarray(result.outputs['output'])
    print(decode_predictions(prediction))
```

Run inference on the model by using gRPC with the following command:

python tensorflow-model-server-infer.py

Your output should look like the following:

```
[[('n02123045', 'tabby', 0.6918919), ('n02127052', 'lynx', 0.12770271), ('n02123159',
 'tiger_cat', 0.08277027), ('n02124075', 'Egyptian_cat', 0.06418919), ('n02128757',
 'snow_leopard', 0.009290541)]]
```

Using MXNet-Neuron and the Amazon Neuron Compiler

The MXNet-Neuron compilation API provides a method to compile a model graph that you can run on an Amazon Inferentia device.

In this example, you use the API to compile a ResNet-50 model and use it to run inference.

For more information about the Neuron SDK, see the Amazon Neuron SDK documentation.

Contents

- Prerequisites
- <u>Activate the Conda Environment</u>
- Resnet50 Compilation
- <u>ResNet50 Inference</u>

Prerequisites

Before using this tutorial, you should have completed the set up steps in <u>Launching a DLAMI</u> <u>Instance with Amazon Neuron</u>. You should also have a familiarity with deep learning and using the DLAMI.

Activate the Conda Environment

Activate the MXNet-Neuron conda environment using the following command:

source activate aws_neuron_mxnet_p36

To exit the current conda environment, run:

```
source deactivate
```

Resnet50 Compilation

Create a Python script called **mxnet_compile_resnet50.py** with the following content. This script uses the MXNet-Neuron compilation Python API to compile a ResNet-50 model.

```
import mxnet as mx
import numpy as np
print("downloading...")
path='http://data.mxnet.io/models/imagenet/'
mx.test_utils.download(path+'resnet/50-layers/resnet-50-0000.params')
mx.test_utils.download(path+'resnet/50-layers/resnet-50-symbol.json')
print("download finished.")
```
```
sym, args, aux = mx.model.load_checkpoint('resnet-50', 0)
print("compile for inferentia using neuron... this will take a few minutes...")
inputs = { "data" : mx.nd.ones([1,3,224,224], name='data', dtype='float32') }
sym, args, aux = mx.contrib.neuron.compile(sym, args, aux, inputs)
print("save compiled model...")
mx.model.save_checkpoint("compiled_resnet50", 0, sym, args, aux)
```

Compile the model using the following command:

```
python mxnet_compile_resnet50.py
```

Compilation will take a few minutes. When compilation has finished, the following files will be in your current directory:

```
resnet-50-0000.params
resnet-50-symbol.json
compiled_resnet50-0000.params
compiled_resnet50-symbol.json
```

ResNet50 Inference

Create a Python script called **mxnet_infer_resnet50.py** with the following content. This script downloads a sample image and uses it to run inference with the compiled model.

```
import mxnet as mx
import numpy as np
path='http://data.mxnet.io/models/imagenet/'
mx.test_utils.download(path+'synset.txt')
fname = mx.test_utils.download('https://raw.githubusercontent.com/awslabs/mxnet-model-
server/master/docs/images/kitten_small.jpg')
img = mx.image.imread(fname)
# convert into format (batch, RGB, width, height)
```

```
img = mx.image.imresize(img, 224, 224)
# resize
img = img.transpose((2, 0, 1))
# Channel first
img = img.expand_dims(axis=0)
# batchify
img = img.astype(dtype='float32')
sym, args, aux = mx.model.load_checkpoint('compiled_resnet50', 0)
softmax = mx.nd.random_normal(shape=(1,))
args['softmax_label'] = softmax
args['data'] = img
# Inferentia context
ctx = mx.neuron()
exe = sym.bind(ctx=ctx, args=args, aux_states=aux, grad_req='null')
with open('synset.txt', 'r') as f:
    labels = [l.rstrip() for l in f]
exe.forward(data=img)
prob = exe.outputs[0].asnumpy()
# print the top-5
prob = np.squeeze(prob)
a = np.argsort(prob)[::-1]
for i in a[0:5]:
    print('probability=%f, class=%s' %(prob[i], labels[i]))
```

Run inference with the compiled model using the following command:

python mxnet_infer_resnet50.py

Your output should look like the following:

```
probability=0.642454, class=n02123045 tabby, tabby cat
probability=0.189407, class=n02123159 tiger cat
probability=0.100798, class=n02124075 Egyptian cat
probability=0.030649, class=n02127052 lynx, catamount
probability=0.016278, class=n02129604 tiger, Panthera tigris
```

Next Step

Using MXNet-Neuron Model Serving

Using MXNet-Neuron Model Serving

In this tutorial, you learn to use a pre-trained MXNet model to perform real-time image classification with Multi Model Server (MMS). MMS is a flexible and easy-to-use tool for serving deep learning models that are trained using any machine learning or deep learning framework. This tutorial includes a compilation step using Amazon Neuron and an implementation of MMS using MXNet.

For more information about the Neuron SDK, see the Amazon Neuron SDK documentation.

Contents

- Prerequisites
- Activate the Conda Environment
- Download the Example Code
- Compile the Model
- Run Inference

Prerequisites

Before using this tutorial, you should have completed the set up steps in <u>Launching a DLAMI</u> <u>Instance with Amazon Neuron</u>. You should also have a familiarity with deep learning and using the DLAMI.

Activate the Conda Environment

Activate the MXNet-Neuron conda environment by using the following command:

```
source activate aws_neuron_mxnet_p36
```

To exit the current conda environment, run:

source deactivate

Download the Example Code

To run this example, download the example code using the following commands:

git clone https://github.com/awslabs/multi-model-server

```
cd multi-model-server/examples/mxnet_vision
```

Compile the Model

Create a Python script called multi-model-server-compile.py with the following content. This script compiles the ResNet50 model to the Inferentia device target.

```
import mxnet as mx
from mxnet.contrib import neuron
import numpy as np
path='http://data.mxnet.io/models/imagenet/'
mx.test_utils.download(path+'resnet/50-layers/resnet-50-0000.params')
mx.test_utils.download(path+'resnet/50-layers/resnet-50-symbol.json')
mx.test_utils.download(path+'synset.txt')
nn_name = "resnet-50"
#Load a model
sym, args, auxs = mx.model.load_checkpoint(nn_name, 0)
#Define compilation parameters# - input shape and dtype
inputs = {'data' : mx.nd.zeros([1,3,224,224], dtype='float32') }
# compile graph to inferentia target
csym, cargs, cauxs = neuron.compile(sym, args, auxs, inputs)
# save compiled model
mx.model.save_checkpoint(nn_name + "_compiled", 0, csym, cargs, cauxs)
```

To compile the model, use the following command:

python multi-model-server-compile.py

Your output should look like the following:

```
...
[21:18:40] src/nnvm/legacy_json_util.cc:209: Loading symbol saved by previous version
v0.8.0. Attempting to upgrade...
[21:18:40] src/nnvm/legacy_json_util.cc:217: Symbol successfully upgraded!
[21:19:00] src/operator/subgraph/build_subgraph.cc:698: start to execute partition
graph.
```

```
[21:19:00] src/nnvm/legacy_json_util.cc:209: Loading symbol saved by previous version
v0.8.0. Attempting to upgrade...
[21:19:00] src/nnvm/legacy_json_util.cc:217: Symbol successfully upgraded!
```

Create a file named signature.json with the following content to configure the input name and shape:

Download the synset.txt file by using the following command. This file is a list of names for ImageNet prediction classes.

```
curl -0 https://s3.amazonaws.com/model-server/model_archive_1.0/examples/
squeezenet_v1.1/synset.txt
```

Create a custom service class following the template in the model_server_template folder. Copy the template into your current working directory by using the following command:

```
cp -r ../model_service_template/* .
```

Edit the mxnet_model_service.py module to replace the mx.cpu() context with the mx.neuron() context as follows. You also need to comment out the unnecessary data copy for model_input because MXNet-Neuron does not support the NDArray and Gluon APIs.

```
...
self.mxnet_ctx = mx.neuron() if gpu_id is None else mx.gpu(gpu_id)
...
#model_input = [item.as_in_context(self.mxnet_ctx) for item in model_input]
```

Package the model with model-archiver using the following commands:

```
cd ~/multi-model-server/examples
model-archiver --force --model-name resnet-50_compiled --model-path mxnet_vision --
handler mxnet_vision_service:handle
```

Run Inference

Start the Multi Model Server and load the model that uses the RESTful API by using the following commands. Ensure that **neuron-rtd** is running with the default settings.

```
cd ~/multi-model-server/
multi-model-server --start --model-store examples > /dev/null # Pipe to log file if you
want to keep a log of MMS
curl -v -X POST "http://localhost:8081/models?
initial_workers=1&max_workers=4&synchronous=true&url=resnet-50_compiled.mar"
sleep 10 # allow sufficient time to load model
```

Run inference using an example image with the following commands:

```
curl -0 https://raw.githubusercontent.com/awslabs/multi-model-server/master/docs/
images/kitten_small.jpg
curl -X POST http://127.0.0.1:8080/predictions/resnet-50_compiled -T kitten_small.jpg
```

Your output should look like the following:

```
Ε
 {
    "probability": 0.6388034820556641,
    "class": "n02123045 tabby, tabby cat"
 },
  ſ
    "probability": 0.16900072991847992,
    "class": "n02123159 tiger cat"
 },
  {
    "probability": 0.12221276015043259,
    "class": "n02124075 Egyptian cat"
 },
 {
    "probability": 0.028706775978207588,
    "class": "n02127052 lynx, catamount"
```

```
},
{
    "probability": 0.01915954425930977,
    "class": "n02129604 tiger, Panthera tigris"
}
]
```

To cleanup after the test, issue a delete command via the RESTful API and stop the model server using the following commands:

```
curl -X DELETE http://127.0.0.1:8081/models/resnet-50_compiled
multi-model-server --stop
```

You should see the following output:

```
{
    "status": "Model \"resnet-50_compiled\" unregistered"
}
Model server stopped.
Found 1 models and 1 NCGs.
Unloading 10001 (MODEL_STATUS_STARTED) :: success
Destroying NCG 1 :: success
```

Using PyTorch-Neuron and the Amazon Neuron Compiler

The PyTorch-Neuron compilation API provides a method to compile a model graph that you can run on an Amazon Inferentia device.

A trained model must be compiled to an Inferentia target before it can be deployed on Inf1 instances. The following tutorial compiles the torchvision ResNet50 model and exports it as a saved TorchScript module. This model is then used to run inference.

For convenience, this tutorial uses an Inf1 instance for both compilation and inference. In practice, you may compile your model using another instance type, such as the c5 instance family. You must then deploy your compiled model to the Inf1 inference server. For more information, see the Amazon Neuron PyTorch SDK Documentation.

Contents

- Prerequisites
- Activate the Conda Environment

- Resnet50 Compilation
- ResNet50 Inference

Prerequisites

Before using this tutorial, you should have completed the set up steps in <u>Launching a DLAMI</u> <u>Instance with Amazon Neuron</u>. You should also have a familiarity with deep learning and using the DLAMI.

Activate the Conda Environment

Activate the PyTorch-Neuron conda environment using the following command:

source activate aws_neuron_pytorch_p36

To exit the current conda environment, run:

source deactivate

Resnet50 Compilation

Create a Python script called **pytorch_trace_resnet50.py** with the following content. This script uses the PyTorch-Neuron compilation Python API to compile a ResNet-50 model.

Note

There is a dependency between versions of torchvision and the torch package that you should be aware of when compiling torchvision models. These dependency rules can be managed through pip. Torchvision==0.6.1 matches the torch==1.5.1 release, while torchvision==0.8.2 matches the torch==1.7.1 release.

```
import torch
import numpy as np
import os
import torch_neuron
from torchvision import models
image = torch.zeros([1, 3, 224, 224], dtype=torch.float32)
```

```
## Load a pretrained ResNet50 model
model = models.resnet50(pretrained=True)
## Tell the model we are using it for evaluation (not training)
model.eval()
model_neuron = torch.neuron.trace(model, example_inputs=[image])
## Export to saved model
model_neuron.save("resnet50_neuron.pt")
```

Run the compilation script.

python pytorch_trace_resnet50.py

Compilation will take a few minutes. When compilation has finished, the compiled model is saved as resnet50_neuron.pt in the local directory.

ResNet50 Inference

Create a Python script called **pytorch_infer_resnet50.py** with the following content. This script downloads a sample image and uses it to run inference with the compiled model.

Fetch labels to output the top classifications

```
request.urlretrieve("https://s3.amazonaws.com/deep-learning-models/image-models/
imagenet_class_index.json","imagenet_class_index.json")
idx2label = []
with open("imagenet_class_index.json", "r") as read_file:
    class_idx = json.load(read_file)
    idx2label = [class_idx[str(k)][1] for k in range(len(class_idx))]
## Import a sample image and normalize it into a tensor
normalize = transforms.Normalize(
    mean=[0.485, 0.456, 0.406],
    std=[0.229, 0.224, 0.225])
eval_dataset = datasets.ImageFolder(
    os.path.dirname("./torch_neuron_test/"),
    transforms.Compose([
    transforms.Resize([224, 224]),
   transforms.ToTensor(),
    normalize,
    ])
)
image, _ = eval_dataset[0]
image = torch.tensor(image.numpy()[np.newaxis, ...])
## Load model
model_neuron = torch.jit.load( 'resnet50_neuron.pt' )
## Predict
results = model_neuron( image )
# Get the top 5 results
top5_idx = results[0].sort()[1][-5:]
# Lookup and print the top 5 labels
top5_labels = [idx2label[idx] for idx in top5_idx]
print("Top 5 labels:\n {}".format(top5_labels) )
```

Run inference with the compiled model using the following command:

```
python pytorch_infer_resnet50.py
```

Your output should look like the following:

```
Top 5 labels:
['tiger', 'lynx', 'tiger_cat', 'Egyptian_cat', 'tabby']
```

The ARM64 DLAMI

Amazon ARM64 GPU DLAMIs are designed to provide high performance and cost efficiency for deep learning workloads. Specifically, the G5g instance type features the Arm64-based <u>Amazon</u> <u>Graviton2 processor</u>, which was built from the ground up by Amazon and optimized for how customers run their workloads in the cloud. Amazon ARM64 GPU DLAMIs are pre-configured with Docker, NVIDIA Docker, NVIDIA Driver, CUDA, CuDNN, NCCL, as well as popular machine learning frameworks such as TensorFlow and PyTorch.

With the G5g instance type, you can take advantage of the price and performance benefits of Graviton2 to deploy GPU-accelerated deep learning models at a significantly lower cost when compared with x86-based instances with GPU acceleration.

Select a ARM64 DLAMI

Launch a G5g instance with the ARM64 DLAMI of your choice.

For step-by-step instructions on launching a DLAMI, see Launching and Configuring a DLAMI.

For a list of the most recent ARM64 DLAMIs, see the Release Notes for DLAMI.

Get Started

The following topics show you how to get started using the ARM64 DLAMI.

Contents

Using the ARM64 GPU PyTorch DLAMI

Using the ARM64 GPU PyTorch DLAMI

The Amazon Deep Learning AMIs is ready to use with Arm64 processor-based GPUs, and comes optimized for PyTorch. The ARM64 GPU PyTorch DLAMI includes a Python environment pre-configured with <u>PyTorch</u>, <u>TorchVision</u>, and <u>TorchServe</u> for deep learning training and inference use cases.

Contents

- Verify PyTorch Python Environment
- Run Training Sample with PyTorch
- Run Inference Sample with PyTorch

Verify PyTorch Python Environment

Connect to your G5g instance and activate the base Conda environment with the following command:

source activate base

Your command prompt should indicate that you are working in the base Conda environment, which contains PyTorch, TorchVision, and other libraries.

(base) \$

Verify the default tool paths of the PyTorch environment:

```
(base) $ which python
(base) $ which pip
(base) $ which conda
(base) $ which mamba
>>> import torch, torchvision
>>> torch.__version__
>>> torchvision.__version__
>>> v = torch.autograd.Variable(torch.randn(10, 3, 224, 224))
>>> v = torch.autograd.Variable(torch.randn(10, 3, 224, 224)).cuda()
>>> assert isinstance(v, torch.Tensor)
```

Run Training Sample with PyTorch

Run a sample MNIST training job:

```
git clone https://github.com/pytorch/examples.git
cd examples/mnist
python main.py
```

Your output should look similar to the following:

• • •		
Train Epoch:	14 [56320/60000 (94%)] Loss: 0.021424	
Train Epoch:	14 [56960/60000 (95%)] Loss: 0.023695	
Train Epoch:	14 [57600/60000 (96%)] Loss: 0.001973	
Train Epoch:	14 [58240/60000 (97%)] Loss: 0.007121	
Train Epoch:	14 [58880/60000 (98%)] Loss: 0.003717	
Train Epoch:	14 [59520/60000 (99%)] Loss: 0.001729	
Test set: Av	rage loss: 0.0275, Accuracy: 9916/10000 ((99%)

Run Inference Sample with PyTorch

Use the following commands to download a pre-trained densenet161 model and run inference using TorchServe:

```
# Set up TorchServe
cd $HOME
git clone https://github.com/pytorch/serve.git
mkdir -p serve/model_store
cd serve
# Download a pre-trained densenet161 model
wget https://download.pytorch.org/models/densenet161-8d451a50.pth >/dev/null
# Save the model using torch-model-archiver
torch-model-archiver --model-name densenet161 \
    --version 1.0 \setminus
    --model-file examples/image_classifier/densenet_161/model.py \
    --serialized-file densenet161-8d451a50.pth \
    --handler image_classifier \
    --extra-files examples/image_classifier/index_to_name.json \
    --export-path model_store
# Start the model server
torchserve --start --no-config-snapshots \
    --model-store model_store \
    --models densenet161=densenet161.mar &> torchserve.log
# Wait for the model server to start
sleep 30
# Run a prediction request
```

```
curl http://127.0.0.1:8080/predictions/densenet161 -T examples/image_classifier/
kitten.jpg
```

Your output should look similar to the following:

```
{
    "tiger_cat": 0.4693363308906555,
    "tabby": 0.4633873701095581,
    "Egyptian_cat": 0.06456123292446136,
    "lynx": 0.0012828150065615773,
    "plastic_bag": 0.00023322898778133094
}
```

Use the following commands to unregister the densenet161 model and stop the server:

```
curl -X DELETE http://localhost:8081/models/densenet161/1.0
torchserve --stop
```

Your output should look similar to the following:

```
{
   "status": "Model \"densenet161\" unregistered"
}
TorchServe has stopped.
```

Inference

This section provides tutorials on how to run inference using the DLAMI's frameworks and tools.

Inference Tools

TensorFlow Serving

Model Serving

The following are model serving options installed on the Deep Learning AMI with Conda. Click on one of the options to learn how to use it.

Topics

TensorFlow Serving

TorchServe

TensorFlow Serving

TensorFlow Serving is a flexible, high-performance serving system for machine learning models.

The tensorflow-serving-api is pre-installed with single framwork DLAMI. To use tensorflow serving, first activate the TensorFlow environment.

```
$ source /opt/tensorflow/bin/activate
```

Then use your preferred text editor to create a script that has the following content. Name it test_train_mnist.py. This script is referenced from <u>TensorFlow Tutorial</u> which will train and evaluate a neural network machine learning model that classifies images.

Now run the script passing the server location and port and the husky photo's filename as the parameters.

```
$ /opt/tensorflow/bin/python3 test_train_mnist.py
```

Be patient, as this script may take a while before providing any output. When the training is complete you should see the following:

```
10000 00:00:1739482012.389276
                                4284 device_compiler.h:188] Compiled cluster using
XLA! This line is logged at most once for the lifetime of the process.
1875/1875 [==================] - 24s 2ms/step - loss: 0.2973 - accuracy:
0.9134
Epoch 2/5
1875/1875 [=================================] - 3s 2ms/step - loss: 0.1422 - accuracy:
0.9582
Epoch 3/5
1875/1875 [==================================] - 3s 1ms/step - loss: 0.1076 - accuracy:
0.9687
Epoch 4/5
1875/1875 [========================] - 3s 2ms/step - loss: 0.0872 - accuracy:
0.9731
Epoch 5/5
1875/1875 [========================] - 3s 1ms/step - loss: 0.0731 - accuracy:
0.9771
313/313 [========================] - 0s 1ms/step - loss: 0.0749 - accuracy:
0.9780
```

More Features and Examples

If you are interested in learning more about TensorFlow Serving, check out the <u>TensorFlow website</u>.

TorchServe

TorchServe is a flexible tool for serving deep learning models that have been exported from PyTorch. TorchServe comes preinstalled with the Deep Learning AMI with Conda.

For more information on using TorchServe, see Model Server for PyTorch Documentation.

Topics

Serve an Image Classification Model on TorchServe

This tutorial shows how to serve an image classification model with TorchServe. It uses a DenseNet-161 model provided by PyTorch. Once the server is running, it listens for prediction requests. When you upload an image, in this case, an image of a kitten, the server returns a prediction of the top 5 matching classes out of the classes that the model was trained on.

To serve an example image classification model on TorchServe

- 1. Connect to an Amazon Elastic Compute Cloud (Amazon EC2) instance with Deep Learning AMI with Conda v34 or later.
- 2. Activate the pytorch_p310 environment.

source activate pytorch_p310

3. Clone the TorchServe repository, then create a directory to store your models.

```
git clone https://github.com/pytorch/serve.git
mkdir model_store
```

4. Archive the model using the model archiver. The extra-files param uses a file from the TorchServe repo, so update the path if necessary. For more information about the model archiver, see Torch Model archiver for TorchServe.

```
wget https://download.pytorch.org/models/densenet161-8d451a50.pth
torch-model-archiver --model-name densenet161 --version 1.0 --model-file ./
serve/examples/image_classifier/densenet_161/model.py --serialized-file
densenet161-8d451a50.pth --export-path model_store --extra-files ./serve/examples/
image_classifier/index_to_name.json --handler image_classifier
```

5. Run TorchServe to start an endpoint. Adding > /dev/null quiets the log output.

```
torchserve --start --ncs --model-store model_store --models densenet161.mar > /dev/
null
```

6. Download an image of a kitten and send it to the TorchServe predict endpoint:

```
curl -0 https://s3.amazonaws.com/model-server/inputs/kitten.jpg
curl http://127.0.0.1:8080/predictions/densenet161 -T kitten.jpg
```

The predict endpoint returns a prediction in JSON similar to the following top five predictions, where the image has a 47% probability of containing an Egyptian cat, followed by a 46% chance it has a tabby cat.

```
{
    "tiger_cat": 0.46933576464653015,
    "tabby": 0.463387668132782,
```

```
"Egyptian_cat": 0.0645613968372345,
"lynx": 0.0012828196631744504,
"plastic_bag": 0.00023323058849200606
}
```

7. When you finish testing, stop the server:

```
torchserve --stop
```

Other Examples

TorchServe has a variety of examples that you can run on your DLAMI instance. You can view them on the TorchServe project repository examples page.

More Info

For more TorchServe documentation, including how to set up TorchServe with Docker and the latest TorchServe features, see the TorchServe project pageon GitHub.

Upgrading Your DLAMI

Here you will find information on upgrading your DLAMI and tips on updating software on your DLAMI.

Always keep your operating system and other installed software up to date by applying patches and updates as soon as they become available.

If you are using Amazon Linux or Ubuntu, when you login to your DLAMI, you are notified if updates are available and see instructions for updating. For further information on Amazon Linux maintenance, see <u>Updating Instance Software</u>. For Ubuntu instances, refer to the official <u>Ubuntu</u> <u>documentation</u>.

On Windows, check Windows Update regularly for software and security updates. If you prefer, have updates applied automatically.

🛕 Important

For information about the Meltdown and Spectre vulnerabilities and how to patch your operating system to address them, see <u>Security Bulletin Amazon-2018-013</u>.

Topics

- Upgrading to a New DLAMI Version
- Tips for Software Updates
- Receive Notifications on New Updates

Upgrading to a New DLAMI Version

DLAMI's system images are updated on a regular basis to take advantage of new deep learning framework releases, CUDA and other software updates, and performance tuning. If you have been using a DLAMI for some time and want to take advantage of an update, you would need to launch a new instance. You would also have to manually transfer any datasets, checkpoints, or other valuable data. Instead, you may use Amazon EBS to retain your data and attach it to a new DLAMI. In this way, you can upgrade often, while minimizing the time it takes to transition your data.

🚯 Note

When attaching and moving Amazon EBS volumes between DLAMIs, you must have both the DLAMIs and the new volume in the same Availability Zone.

- 1. Use the Amazon EC2console to create a new Amazon EBS volume. For detailed directions, see Creating an Amazon EBS Volume.
- 2. Attach your newly created Amazon EBS volume to your existing DLAMI. For detailed directions, see Attaching an Amazon EBS Volume.
- 3. Transfer your data, such as datasets, checkpoints, and configuration files.
- 4. Launch a DLAMI. For detailed directions, see <u>Setting up a DLAMI instance</u>.
- 5. Detach the Amazon EBS volume from your old DLAMI. For detailed directions, see <u>Detaching</u> an Amazon EBS Volume.
- 6. Attach the Amazon EBS volume to your new DLAMI. Follow the instructions from the Step 2 to attach the volume.
- After you verify that your data is available on your new DLAMI, stop and terminate your old DLAMI. For detailed clean-up instructions, see <u>Cleaning up a DLAMI instance</u>.

Tips for Software Updates

From time to time, you may want to manually update software on your DLAMI. It is generally recommended that you use pip to update Python packages. You should also use pip to update packages within a Conda environment on the Deep Learning AMI with Conda. Refer to the particular framework's or software's website for upgrading and installation instructions.

🚺 Note

We cannot guarantee that a package update will be successful. Attempting to update a package in an environment with incompatible dependencies can result in a failure. In such a case, you should contact the library maintainer to see if it is possible to update the package dependencies. Alternatively, you can attempt to modify the environment in such a way that allows the update. However, this modification will likely mean removing or updating existing packages, which means that we can no longer guarantee stability of this environment. The Amazon Deep Learning AMIs comes with many Conda environments and many packages preinstalled. Due to the number of packages preinstalled, finding a set of packages that are guaranteed to be compatible is difficult. You may see a warning "The environment is inconsistent, please check the package plan carefully". DLAMI ensures that all the DLAMI-provided environments are correct, but cannot guarantee that any user installed packages will function correctly.

Receive Notifications on New Updates

🚯 Note

Amazon Deep Learning AMIs have a weekly release cadence for security patches. Release notifications will be sent for these incremental security patches though they may not be included in official release notes.

You can receive notifications whenever a new DLAMI is released. Notifications are published with <u>Amazon SNS</u> using the following topic.

arn:aws:sns:us-west-2:767397762724:dlami-updates

Messages are posted here when a new DLAMI is published. The version, metadata, and regional AMI ID's of the AMI will be included in the message.

These messages can be received using several different methods. We recommend that you use the following method.

- 1. Open the Amazon SNS console.
- 2. In the navigation bar, change the Amazon Region to **US West (Oregon)**, if necessary. You must select the region where the SNS notification that you're subscribing to was created.
- 3. In the navigation pane, choose **Subscriptions**, **Create subscription**.
- 4. For the **Create subscription** dialog box, do the following:
 - a. For **Topic ARN**, copy and paste the following Amazon Resource Name (ARN): arn:aws:sns:us-west-2:767397762724:dlami-updates
 - b. For Protocol, choose one from [Amazon SQS, Amazon Lamda, Email, Email-JSON]
 - c. For **Endpoint**, enter the email address or **Amazon Resource Name (ARN)** of resource that you will use to receive the notifications.

d. Choose Create subscription.

5. You receive a confirmation email with the subject line *Amazon Notification - Subscription Confirmation*. Open the email and choose **Confirm subscription** to complete your subscription.

Security in Amazon Deep Learning AMIs

Cloud security at Amazon is the highest priority. As an Amazon customer, you benefit from data centers and network architectures that are built to meet the requirements of the most security-sensitive organizations.

Security is a shared responsibility between Amazon and you. The <u>shared responsibility model</u> describes this as security *of* the cloud and security *in* the cloud:

- Security of the cloud Amazon is responsible for protecting the infrastructure that runs Amazon Web Services services in the Amazon Web Services Cloud. Amazon also provides you with services that you can use securely. Third-party auditors regularly test and verify the effectiveness of our security as part of the <u>Amazon Compliance Programs</u>. To learn about the compliance programs that apply to Amazon Deep Learning AMIs, see <u>Amazon Services in Scope</u> by Compliance Program.
- Security in the cloud Your responsibility is determined by the Amazon Web Services service that you use. You are also responsible for other factors including the sensitivity of your data, your company's requirements, and applicable laws and regulations.

This documentation helps you understand how to apply the shared responsibility model when using DLAMI. The following topics show you how to configure DLAMI to meet your security and compliance objectives. You also learn how to use other Amazon Web Services services that help you to monitor and secure your DLAMI resources.

For more information, see <u>Security in Amazon EC2</u> in the Amazon EC2 User Guide.

Topics

- Data protection in Amazon Deep Learning AMIs
- Identity and access management for Amazon Deep Learning AMIs
- <u>Compliance validation for Amazon Deep Learning AMIs</u>
- <u>Resilience in Amazon Deep Learning AMIs</u>
- Infrastructure security in Amazon Deep Learning AMIs
- Monitoring Amazon Deep Learning AMIs instances

Data protection in Amazon Deep Learning AMIs

The Amazon <u>shared responsibility model</u> applies to data protection in Amazon Deep Learning AMIs. As described in this model, Amazon is responsible for protecting the global infrastructure that runs all of the Amazon Web Services Cloud. You are responsible for maintaining control over your content that is hosted on this infrastructure. You are also responsible for the security configuration and management tasks for the Amazon Web Services services that you use. For more information about data privacy, see the Data Privacy FAQ.

For data protection purposes, we recommend that you protect Amazon Web Services account credentials and set up individual users with Amazon IAM Identity Center or Amazon Identity and Access Management (IAM). That way, each user is given only the permissions necessary to fulfill their job duties. We also recommend that you secure your data in the following ways:

- Use multi-factor authentication (MFA) with each account.
- Use SSL/TLS to communicate with Amazon resources. We require TLS 1.2 and recommend TLS 1.3.
- Set up API and user activity logging with Amazon CloudTrail. For information about using CloudTrail trails to capture Amazon activities, see <u>Working with CloudTrail trails</u> in the Amazon CloudTrail User Guide.
- Use Amazon encryption solutions, along with all default security controls within Amazon Web Services services.
- Use advanced managed security services such as Amazon Macie, which assists in discovering and securing sensitive data that is stored in Amazon S3.
- If you require FIPS 140-3 validated cryptographic modules when accessing Amazon through a command line interface or an API, use a FIPS endpoint. For more information about the available FIPS endpoints, see Federal Information Processing Standard (FIPS) 140-3.

We strongly recommend that you never put confidential or sensitive information, such as your customers' email addresses, into tags or free-form text fields such as a **Name** field. This includes when you work with DLAMI or other Amazon Web Services services using the console, API, Amazon CLI, or Amazon SDKs. Any data that you enter into tags or free-form text fields used for names may be used for billing or diagnostic logs. If you provide a URL to an external server, we strongly recommend that you do not include credentials information in the URL to validate your request to that server.

Identity and access management for Amazon Deep Learning AMIs

Amazon Identity and Access Management (IAM) is an Amazon Web Services service that helps an administrator securely control access to Amazon resources. IAM administrators control who can be *authenticated* (signed in) and *authorized* (have permissions) to use DLAMI resources. IAM is an Amazon Web Services service that you can use with no additional charge.

For more information about identity and access management, see <u>Identity and access management</u> for Amazon EC2.

Topics

- Authenticating with identities
- Managing access using policies
- IAM with Amazon EMR

Authenticating with identities

Authentication is how you sign in to Amazon using your identity credentials. You must be *authenticated* (signed in to Amazon) as the Amazon Web Services account root user, as an IAM user, or by assuming an IAM role.

If you access Amazon programmatically, Amazon provides a software development kit (SDK) and a command line interface (CLI) to cryptographically sign your requests by using your credentials. If you don't use Amazon tools, you must sign requests yourself. For more information about using the recommended method to sign requests yourself, see <u>Amazon Signature Version 4 for API requests</u> in the *IAM User Guide*.

Regardless of the authentication method that you use, you might be required to provide additional security information. For example, Amazon recommends that you use multi-factor authentication (MFA) to increase the security of your account. To learn more, see <u>Amazon Multi-factor authentication in IAM</u> in the *IAM User Guide*.

Amazon Web Services account root user

When you create an Amazon Web Services account, you begin with one sign-in identity that has complete access to all Amazon Web Services services and resources in the account. This identity

is called the Amazon Web Services account *root user* and is accessed by signing in with the email address and password that you used to create the account. We strongly recommend that you don't use the root user for your everyday tasks. Safeguard your root user credentials and use them to perform the tasks that only the root user can perform. For the complete list of tasks that require you to sign in as the root user, see Tasks that require root user credentials in the *IAM User Guide*.

IAM users and groups

An <u>IAM user</u> is an identity within your Amazon Web Services account that has specific permissions for a single person or application. Where possible, we recommend relying on temporary credentials instead of creating IAM users who have long-term credentials such as passwords and access keys. However, if you have specific use cases that require long-term credentials with IAM users, we recommend that you rotate access keys. For more information, see <u>Rotate access keys regularly for</u> use cases that require long-term credentials in the *IAM User Guide*.

An <u>IAM group</u> is an identity that specifies a collection of IAM users. You can't sign in as a group. You can use groups to specify permissions for multiple users at a time. Groups make permissions easier to manage for large sets of users. For example, you could have a group named *IAMAdmins* and give that group permissions to administer IAM resources.

Users are different from roles. A user is uniquely associated with one person or application, but a role is intended to be assumable by anyone who needs it. Users have permanent long-term credentials, but roles provide temporary credentials. To learn more, see <u>Use cases for IAM users</u> in the *IAM User Guide*.

IAM roles

An <u>IAM role</u> is an identity within your Amazon Web Services account that has specific permissions. It is similar to an IAM user, but is not associated with a specific person. To temporarily assume an IAM role in the Amazon Web Services Management Console, you can <u>switch from a user to an IAM</u> <u>role (console)</u>. You can assume a role by calling an Amazon CLI or Amazon API operation or by using a custom URL. For more information about methods for using roles, see <u>Methods to assume a</u> role in the *IAM User Guide*.

IAM roles with temporary credentials are useful in the following situations:

• **Federated user access** – To assign permissions to a federated identity, you create a role and define permissions for the role. When a federated identity authenticates, the identity is associated with the role and is granted the permissions that are defined by the role. For information about roles for federation, see <u>Create a role for a third-party identity provider</u> (federation) in the *IAM User Guide*.

- **Temporary IAM user permissions** An IAM user or role can assume an IAM role to temporarily take on different permissions for a specific task.
- Cross-account access You can use an IAM role to allow someone (a trusted principal) in a different account to access resources in your account. Roles are the primary way to grant cross-account access. However, with some Amazon Web Services services, you can attach a policy directly to a resource (instead of using a role as a proxy). To learn the difference between roles and resource-based policies for cross-account access, see Cross account resource access in IAM in the IAM User Guide.
- **Cross-service access** Some Amazon Web Services services use features in other Amazon Web Services services. For example, when you make a call in a service, it's common for that service to run applications in Amazon EC2 or store objects in Amazon S3. A service might do this using the calling principal's permissions, using a service role, or using a service-linked role.
 - Forward access sessions (FAS) When you use an IAM user or role to perform actions in Amazon, you are considered a principal. When you use some services, you might perform an action that then initiates another action in a different service. FAS uses the permissions of the principal calling an Amazon Web Services service, combined with the requesting Amazon Web Services service to make requests to downstream services. FAS requests are only made when a service receives a request that requires interactions with other Amazon Web Services services or resources to complete. In this case, you must have permissions to perform both actions. For policy details when making FAS requests, see Forward access sessions.
 - Service role A service role is an <u>IAM role</u> that a service assumes to perform actions on your behalf. An IAM administrator can create, modify, and delete a service role from within IAM. For more information, see <u>Create a role to delegate permissions to an Amazon Web Services</u> service in the *IAM User Guide*.
 - Service-linked role A service-linked role is a type of service role that is linked to an Amazon Web Services service. The service can assume the role to perform an action on your behalf. Service-linked roles appear in your Amazon Web Services account and are owned by the service. An IAM administrator can view, but not edit the permissions for service-linked roles.
- Applications running on Amazon EC2 You can use an IAM role to manage temporary credentials for applications that are running on an EC2 instance and making Amazon CLI or Amazon API requests. This is preferable to storing access keys within the EC2 instance. To assign an Amazon role to an EC2 instance and make it available to all of its applications, you create an instance profile that is attached to the instance. An instance profile contains the role and

enables programs that are running on the EC2 instance to get temporary credentials. For more information, see <u>Use an IAM role to grant permissions to applications running on Amazon EC2</u> instances in the *IAM User Guide*.

Managing access using policies

You control access in Amazon by creating policies and attaching them to Amazon identities or resources. A policy is an object in Amazon that, when associated with an identity or resource, defines their permissions. Amazon evaluates these policies when a principal (user, root user, or role session) makes a request. Permissions in the policies determine whether the request is allowed or denied. Most policies are stored in Amazon as JSON documents. For more information about the structure and contents of JSON policy documents, see <u>Overview of JSON policies</u> in the *IAM User Guide*.

Administrators can use Amazon JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

By default, users and roles have no permissions. To grant users permission to perform actions on the resources that they need, an IAM administrator can create IAM policies. The administrator can then add the IAM policies to roles, and users can assume the roles.

IAM policies define permissions for an action regardless of the method that you use to perform the operation. For example, suppose that you have a policy that allows the iam:GetRole action. A user with that policy can get role information from the Amazon Web Services Management Console, the Amazon CLI, or the Amazon API.

Identity-based policies

Identity-based policies are JSON permissions policy documents that you can attach to an identity, such as an IAM user, group of users, or role. These policies control what actions users and roles can perform, on which resources, and under what conditions. To learn how to create an identity-based policy, see <u>Define custom IAM permissions with customer managed policies</u> in the *IAM User Guide*.

Identity-based policies can be further categorized as *inline policies* or *managed policies*. Inline policies are embedded directly into a single user, group, or role. Managed policies are standalone policies that you can attach to multiple users, groups, and roles in your Amazon Web Services account. Managed policies include Amazon managed policies and customer managed policies. To learn how to choose between a managed policy or an inline policy, see <u>Choose between managed</u> policies and inline policies in the *IAM User Guide*.

Resource-based policies

Resource-based policies are JSON policy documents that you attach to a resource. Examples of resource-based policies are IAM *role trust policies* and Amazon S3 *bucket policies*. In services that support resource-based policies, service administrators can use them to control access to a specific resource. For the resource where the policy is attached, the policy defines what actions a specified principal can perform on that resource and under what conditions. You must <u>specify a principal</u> in a resource-based policy. Principals can include accounts, users, roles, federated users, or Amazon Web Services services.

Resource-based policies are inline policies that are located in that service. You can't use Amazon managed policies from IAM in a resource-based policy.

Access control lists (ACLs)

Access control lists (ACLs) control which principals (account members, users, or roles) have permissions to access a resource. ACLs are similar to resource-based policies, although they do not use the JSON policy document format.

Amazon S3, Amazon WAF, and Amazon VPC are examples of services that support ACLs. To learn more about ACLs, see <u>Access control list (ACL) overview</u> in the *Amazon Simple Storage Service Developer Guide*.

Other policy types

Amazon supports additional, less-common policy types. These policy types can set the maximum permissions granted to you by the more common policy types.

- Permissions boundaries A permissions boundary is an advanced feature in which you set the maximum permissions that an identity-based policy can grant to an IAM entity (IAM user or role). You can set a permissions boundary for an entity. The resulting permissions are the intersection of an entity's identity-based policies and its permissions boundaries. Resource-based policies that specify the user or role in the Principal field are not limited by the permissions boundary. An explicit deny in any of these policies overrides the allow. For more information about permissions boundaries, see Permissions boundaries for IAM entities in the IAM User Guide.
- Service control policies (SCPs) SCPs are JSON policies that specify the maximum permissions for an organization or organizational unit (OU) in Amazon Organizations. Amazon Organizations is a service for grouping and centrally managing multiple Amazon Web Services accounts that your business owns. If you enable all features in an organization, then you can apply service

control policies (SCPs) to any or all of your accounts. The SCP limits permissions for entities in member accounts, including each Amazon Web Services account root user. For more information about Organizations and SCPs, see <u>Service control policies</u> in the *Amazon Organizations User Guide*.

- Resource control policies (RCPs) RCPs are JSON policies that you can use to set the maximum available permissions for resources in your accounts without updating the IAM policies attached to each resource that you own. The RCP limits permissions for resources in member accounts and can impact the effective permissions for identities, including the Amazon Web Services account root user, regardless of whether they belong to your organization. For more information about Organizations and RCPs, including a list of Amazon Web Services services that support RCPs, see <u>Resource control policies (RCPs)</u> in the *Amazon Organizations User Guide*.
- Session policies Session policies are advanced policies that you pass as a parameter when you
 programmatically create a temporary session for a role or federated user. The resulting session's
 permissions are the intersection of the user or role's identity-based policies and the session
 policies. Permissions can also come from a resource-based policy. An explicit deny in any of these
 policies overrides the allow. For more information, see <u>Session policies</u> in the *IAM User Guide*.

Multiple policy types

When multiple types of policies apply to a request, the resulting permissions are more complicated to understand. To learn how Amazon determines whether to allow a request when multiple policy types are involved, see <u>Policy evaluation logic</u> in the *IAM User Guide*.

IAM with Amazon EMR

You can use IAM with Amazon EMR to define users, Amazon resources, groups, roles, and policies. You can also control which Amazon Web Services services these users and roles can access.

For more information about using IAM with Amazon EMR, see <u>Amazon Identity and Access</u> Management for Amazon EMR.

Compliance validation for Amazon Deep Learning AMIs

Third-party auditors assess the security and compliance of Amazon Deep Learning AMIs as part of multiple Amazon compliance programs. For information about the supported compliance programs, see Compliance validation for Amazon EC2. For a list of Amazon Web Services services in scope of specific compliance programs, see <u>Amazon</u> <u>Services in Scope by Compliance Program</u>. For general information, see <u>Amazon Compliance</u> <u>Programs</u>.

You can download third-party audit reports using Amazon Artifact. For more information, see Downloading reports in Amazon Artifact.

Your compliance responsibility when using DLAMI is determined by the sensitivity of your data, your company's compliance objectives, and applicable laws and regulations. Amazon provides the following resources to help with compliance:

- <u>Security and Compliance Quick Start GuidesSecurity and Compliance Quick Start Guides</u> These deployment guides discuss architectural considerations and provide steps for deploying securityand compliance-focused baseline environments on Amazon.
- <u>Amazon Compliance Resources</u> This collection of workbooks and guides might apply to your industry and location.
- <u>Evaluating Resources with Amazon Config Rules</u> in the Amazon Config Developer Guide The Amazon Config service assesses how well your resource configurations comply with internal practices, industry guidelines, and regulations.
- <u>Amazon Security Hub</u> This Amazon Web Services service provides a comprehensive view of your security state within Amazon. Security Hub uses security controls to evaluate your Amazon resources and to check your compliance against security industry standards and best practices.

Resilience in Amazon Deep Learning AMIs

The Amazon global infrastructure is built around Amazon Web Services Regions and Availability Zones. Amazon Web Services Regions provide multiple physically separated and isolated Availability Zones, which are connected with low-latency, high-throughput, and highly redundant networking. With Availability Zones, you can design and operate applications and databases that automatically fail over between zones without interruption. Availability Zones are more highly available, fault tolerant, and scalable than traditional single or multiple data center infrastructures.

For more information about Amazon Web Services Regions and Availability Zones, see <u>Amazon</u> <u>Global Infrastructure</u>.

For information about Amazon EC2 features to help support your data resiliency and backup needs, see <u>Resilience in Amazon EC2</u> in the *Amazon EC2 User Guide*.

Infrastructure security in Amazon Deep Learning AMIs

The infrastructure security of Amazon Deep Learning AMIs is backed by Amazon EC2. For more information, see Infrastructure security in Amazon EC2 in the Amazon EC2 User Guide.

Monitoring Amazon Deep Learning AMIs instances

Monitoring is an important part of maintaining the reliability, availability, and performance of your Amazon Deep Learning AMIs instance and your other Amazon solutions. Your DLAMI instance comes with several GPU monitoring tools, including a utility that reports GPU usage statistics to Amazon CloudWatch. For more information, see <u>GPU Monitoring and Optimization</u>, and see <u>Monitor Amazon EC2 resources</u> in the *Amazon EC2 User Guide*.

Opting out of usage tracking for DLAMI instances

The following Amazon Deep Learning AMIs operating system distributions include code that allows Amazon to collect instance type, instance ID, DLAMI type, and OS information.

🚯 Note

Amazon doesn't collect or retain any other information about the DLAMI, such as the commands that you use within the DLAMI.

- Amazon Linux 2
- Amazon Linux 2023
- Ubuntu 20.04
- Ubuntu 22.04

To opt out of usage tracking

If you choose, you can opt out of usage tracking for a new DLAMI instance. To opt out, you must add a tag to your Amazon EC2 instance during launch. The tag should use the key OPT_OUT_TRACKING with the associated value set to true. For more information, see <u>Tag your</u> Amazon EC2 resources in the Amazon EC2 User Guide.

DLAMI Support Policy

Here you can find details of the support policy for Amazon Deep Learning AMIs (DLAMI).

For a list of the DLAMI frameworks and operating system that Amazon currently supports, see the <u>DLAMI Support Policy</u> page. The following terminology applies to all DLAMIs mentioned in the Support policy page and this page:

- **Current version** specifies the framework version in the format *x.y.z.* In this format, *x* refers to the major version, *y* refers to the minor version, and *z* refers to the patch version. For example, for TensorFlow 2.10.1, the major version is 2, the minor version is 10, and the patch version is 1.
- End of patch specifies how long Amazon supports a particular framework or operating system version.

For detailed information about specific DLAMIs, see Deep Learning AMIs Release Notes.

DLAMI Support FAQs

- What framework versions get security patches?
- Which operating system get security patches?
- What images does Amazon publish when new framework versions are released?
- What images get new SageMaker AI/Amazon features?
- How is current version defined in the Supported Frameworks table?
- What if I am running a version that is not in the Supported table?
- Do DLAMIs support previous patch versions of a Framework Version?
- How can I find the latest patched image for a supported framework version?
- How frequently are new images released?
- Will my instance be patched in place while my workload is running?
- What happens when a new patched or updated framework version is available?
- Are dependencies updated without changing the framework version?
- When does active support for my framework version end?
- Will images with framework versions that are no longer actively maintained be patched?
- How do I use an older framework version?

- How do I stay up-to-date with support changes in frameworks and their versions?
- Do I need a commercial license to use the Anaconda Repository?

What framework versions get security patches?

If the framework version is under **Supported Framework Versions** in the <u>Amazon Deep Learning</u> AMIs Support Policy table, it gets security patches.

Which operating system get security patches?

If the operating system is listed under **Supported Operating System Versions** in the <u>Amazon Deep</u> <u>Learning AMIs Support Policy table</u>, it gets security patches.

What images does Amazon publish when new framework versions are released?

We publish new DLAMIs soon after new versions of TensorFlow and PyTorch are released. This includes major versions, major-minor versions, and major-minor-patch versions of frameworks. We also update images when new versions of drivers and libraries become available. For more information on image maintenance, see <u>When does active support for my framework version end?</u>

What images get new SageMaker AI/Amazon features?

New features typically release in the latest version of DLAMIs for PyTorch and TensorFlow. Refer to the release notes for a specific image for details on new SageMaker AI or Amazon features. For a list of available DLAMIs, see <u>Release Notes for DLAMI</u>. For more information on image maintenance, see <u>When does active support for my framework version end?</u>

How is current version defined in the Supported Frameworks table?

The current version in the <u>Amazon Deep Learning AMIs Support Policy table</u> refers to the newest framework version that Amazon makes available on GitHub. Each latest release includes updates to the drivers, libraries, and relevant packages in the DLAMI. For information on image maintenance, see <u>When does active support for my framework version end?</u>

What if I am running a version that is not in the Supported table?

If you are running a version that is not in the <u>Amazon Deep Learning AMIs Support Policy table</u>, you may not have the most updated drivers, libraries, and relevant packages. For a more up-to-

date version, we recommend that you upgrade to one of the supported frameworks or operating system available using the latest DLAMI of your choice. For a list of available DLAMIs, see <u>Release</u> Notes for DLAMI.

Do DLAMIs support previous patch versions of a Framework Version?

No. We support the latest patch version of each framework's latest major version released 365 days from its initial GitHub release as stated in the <u>Amazon Deep Learning AMIs Support Policy table</u>. For more information, see <u>What if I am running a version that is not in the Supported table?</u>

How can I find the latest patched image for a supported framework version?

To use a DLAMI with the latest framework version, you can use Amazon CLI or SSM parameters to retrieve the <u>DLAMI ID</u> and use it to launch the DLAMI using the <u>EC2 Console</u>. For sample Amazon CLI or SSM parameter commands to retrieve the Amazon Deep Learning AMIs ID, refer to the DLAMI release notes page <u>single-framework DLAMI release notes</u>. The framework version that you choose must be listed under **Supported Framework Versions** in the <u>Amazon Deep Learning AMIs</u> <u>Support Policy table</u>.

How frequently are new images released?

Providing updated patch versions is our highest priority. We routinely create patched images at the earliest opportunity. We monitor for newly patched framework versions (ex. TensorFlow 2.9 to TensorFlow 2.9.1) and new minor release versions (ex. TensorFlow 2.9 to TensorFlow 2.10) and make them available at the earliest opportunity. When an existing version of TensorFlow is released with a new version of CUDA, we release a new DLAMI for that version of TensorFlow with support for the new CUDA version.

Will my instance be patched in place while my workload is running?

No. Patch updates for DLAMI are not "in-place" updates.

You must turn on a new EC2 instance, migrate your workloads and scripts, and then turn off your previous instance.

What happens when a new patched or updated framework version is available?

To be notified of changes in DLAMI, please subscribe to the notifications for the relevant DLAMI, see <u>Receive Notifications on New Updates</u>.

Are dependencies updated without changing the framework version?

We update dependencies without changing the framework version. However, if a dependency update causes an incompatibility, we create an image with a different version. Be sure to check the <u>Release Notes for DLAMI</u> for updated dependency information.

When does active support for my framework version end?

DLAMI images are immutable. Once they are created they do not change. There are four main reasons why active support for a framework version ends:

- Framework version (patch) upgrades
- <u>Amazon security patches</u>
- End of patch date (Aging out)
- Dependency end-of-support

🚯 Note

Due to the frequency of version patch upgrades and security patches, we recommend checking the release notes page for your DLAMI often, and upgrading when changes are made.

Framework version (patch) upgrades

If you have a DLAMI workload based on TensorFlow 2.7.0 and TensorFlow releases version 2.7.1 on GitHub, then Amazon releases a new DLAMI with TensorFlow 2.7.1. The previous images with 2.7.0 are no longer actively maintained once the new image with TensorFlow 2.7.1 is released. The DLAMI with TensorFlow 2.7.0 does not receive further patches. The DLAMI release notes page for TensorFlow 2.7 is then updated with the latest information. There is no individual release note page for each minor patch.
New DLAMIs created due to patch upgrades are designated with a new AMI ID.

Amazon security patches

If you have a workload based on an image with TensorFlow 2.7.0 and Amazon makes a security patch, then a new version of the DLAMI is released for TensorFlow 2.7.0. The previous version of the images with TensorFlow 2.7.0 is no longer actively maintained. For more information, see <u>Will</u> <u>my instance be patched in place while my workload is running?</u> For steps on finding the latest DLAMI, see How can I find the latest patched image for a supported framework version?

New DLAMIs created due to patch upgrades are designated with a new AMI ID.

End of patch date (Aging out)

DLAMIs hit their end of patch date 365 days after the GitHub release date.

For <u>multi-framework DLAMIs</u>, when one of the framework versions is updated, a new DLAMI with the updated version is required. The DLAMI with the old framework version is no longer actively maintained.

<u> Important</u>

We make an exception when there is a major framework update. For example. if TensorFlow 1.15 updates to TensorFlow 2.0, then we continue to support the most recent version of TensorFlow 1.15 for a period of two years from the date of the GitHub release or six months after the origin framework maintenance team drops support, whichever date is earlier.

Dependency end-of-support

If you are running a workload on a TensorFlow 2.7.0 DLAMI image with Python 3.6 and that version of Python is marked for end-of-support, then all DLAMI images based on Python 3.6 will no longer be actively maintained. Similarly, if an OS version like Ubuntu 16.04 is marked for end-of-support, then all DLAMI images that are dependent on Ubuntu 16.04 will no longer be actively maintained.

Will images with framework versions that are no longer actively maintained be patched?

No. Images that are no longer actively maintained will not have new releases.

How do I use an older framework version?

To use a DLAMI with an older framework version, retrieve the <u>DLAMI ID</u> and use it to launch the DLAMI using the <u>EC2 Console</u>. For Amazon CLI commands to retrieve the AMI ID, refer to the release notes page in the <u>single-framework DLAMI release notes</u>.

How do I stay up-to-date with support changes in frameworks and their versions?

Stay up-to-date with DLAMI frameworks and versions using the <u>Amazon Deep Learning AMIs</u> Framework Support Policy table, the <u>DLAMI release notes</u>.

Do I need a commercial license to use the Anaconda Repository?

Anaconda shifted to a commercial licensing model for certain users. Actively maintained DLAMIs have been migrated to the publicly available open-source version of Conda (<u>conda-forge</u>) from the Anaconda channel.

DLAMI Support Policy Table

For more details see the <u>Support Policy</u>.

Supported Framework Versions

Framework	Current version	CUDA version	GitHub GA	End of patch
PyTorch	2.7.0	12.8	2025-04-23	2026-04-23
PyTorch	2.6.0	12.6	2025-01-29	2026-01-29
PyTorch	2.5.1	12.4	2024-11-24	2025-11-24
PyTorch	2.4.1	12.4	2024-07-24	2025-07-24
TensorFlow	2.18.0	12.5	2024-10-24	2025-10-24
TensorFlow	2.17.0	12.3	2024-11-07	2025-11-07

Supported Operating System Versions

Operating System	End of patch
Amazon Linux 2023	2029-06-30
Amazon Linux 2	2026-06-30
Ubuntu 24.04	2029-04-30
Ubuntu 22.04	2027-04-30

Unsupported Framework Versions

Versions listed in this table will appear for 2 years past their support date.

Framework	Current version	CUDA version	GitHub GA	End of patch
PyTorch	2.3.0	12.1	2024-04-24	2025-04-24
PyTorch	2.2.0	12.1	2024-01-30	2025-01-30
PyTorch	1.13.1	11.7	2022-10-28	2024-10-28
PyTorch	2.1.0	12.1	2023-10-04	2024-10-04
PyTorch	2.0.0	12.1	2023-03-15	2024-03-15
PyTorch	1.12.1	11.6	2022-07-01	2023-07-01
PyTorch	1.11.0	11.5	2022-03-10	2023-03-10
TensorFlow	2.16.0	12.3	2024-03-07	2025-03-07
TensorFlow	2.15.0	12.2	2023-11-14	2024-11-14
TensorFlow	2.13.0	11.8	2023-07-19	2024-07-19
TensorFlow	2.12.0	11.8	2023-03-23	2024-03-23
TensorFlow	2.11.0	11.2	2022-11-18	2023-11-18
TensorFlow	2.10.1	11.2	2022-09-06	2023-09-06
TensorFlow	2.9.3	11.2	2022-05-17	2023-05-17

Unsupported Operating System Versions

Operating System	End of patch
Ubuntu 20.04	2025-05-31
Ubuntu 18.04	2023-05-31

Unsupported DLAMI Release Notes Archive

Base

GPU

- Amazon Deep Learning Base AMI (Ubuntu 20.04)
- Amazon Deep Learning Base AMI (Ubuntu 18.04)

Single Framework

PyTorch Specific AMI

GPU

- Amazon Deep Learning AMI GPU PyTorch 2.3 (Ubuntu 20.04)
- Amazon Deep Learning AMI GPU PyTorch 2.3 (Amazon Linux 2)
- Amazon Deep Learning ARM64 AMI GPU PyTorch 2.3 (Ubuntu 22.04)
- Amazon Deep Learning AMI GPU PyTorch 2.2 (Amazon Linux 2)
- Amazon Deep Learning AMI GPU PyTorch 2.2 (Ubuntu 20.04)
- Amazon Deep Learning AMI GPU PyTorch 2.1 (Ubuntu 20.04)
- Amazon Deep Learning AMI GPU PyTorch 2.0 (Amazon Linux 2)
- Amazon Deep Learning AMI GPU PyTorch 2.0 (Ubuntu 20.04)
- Amazon Deep Learning AMI GPU PyTorch 1.13 (Amazon Linux 2)
- Amazon Deep Learning AMI GPU PyTorch 1.13 (Ubuntu 20.04)
- Amazon Deep Learning AMI GPU PyTorch 1.12 (Amazon Linux 2)
- Amazon Deep Learning AMI GPU PyTorch 1.12 (Ubuntu 20.04)
- Amazon Deep Learning AMI GPU PyTorch 1.11 (Amazon Linux 2)
- Amazon Deep Learning AMI GPU PyTorch 1.11 (Ubuntu 20.04)
- Amazon Deep Learning AMI GPU PyTorch 1.10 (Amazon Linux 2)
- Amazon Deep Learning AMI GPU PyTorch 1.10 (Ubuntu 20.04)
- Amazon Deep Learning AMI Graviton GPU PyTorch 1.10 (Ubuntu 20.04)
- Amazon Deep Learning AMI GPU PyTorch 1.10 (Ubuntu 18.04)

- Amazon Deep Learning AMI GPU PyTorch 1.9 (Amazon Linux 2)
- Amazon Deep Learning AMI GPU PyTorch 1.9 (Ubuntu 20.04)
- Amazon Deep Learning AMI GPU PyTorch 1.9 (Ubuntu 18.04)

TensorFlow Specific AMI

GPU

- Amazon Deep Learning AMI GPU TensorFlow 2.16 (Amazon Linux 2)
- Amazon Deep Learning AMI GPU TensorFlow 2.16 (Ubuntu 20.04)
- Amazon Deep Learning AMI GPU TensorFlow 2.15 (Amazon Linux 2)
- Amazon Deep Learning AMI GPU TensorFlow 2.15 (Ubuntu 20.04)
- Amazon Deep Learning AMI GPU TensorFlow 2.13 (Amazon Linux 2)
- Amazon Deep Learning AMI GPU TensorFlow 2.13 (Ubuntu 20.04)
- Amazon Deep Learning AMI GPU TensorFlow 2.12 (Amazon Linux 2)
- Amazon Deep Learning AMI GPU TensorFlow 2.12 (Ubuntu 20.04)
- <u>Amazon Deep Learning AMI GPU TensorFlow 2.11 (Amazon Linux 2)</u>
- Amazon Deep Learning AMI GPU TensorFlow 2.11 (Ubuntu 20.04)
- Amazon Deep Learning AMI GPU TensorFlow 2.10 (Amazon Linux 2)
- Amazon Deep Learning AMI GPU TensorFlow 2.10 (Ubuntu 20.04)
- Amazon Deep Learning AMI GPU TensorFlow 2.9 (Amazon Linux 2)
- Amazon Deep Learning AMI GPU TensorFlow 2.9 (Ubuntu 20.04)
- Amazon Deep Learning AMI GPU TensorFlow 2.8 (Amazon Linux 2)
- Amazon Deep Learning AMI GPU TensorFlow 2.8 (Ubuntu 20.04)
- Amazon Deep Learning AMI GPU TensorFlow 2.7 (Amazon Linux 2)
- Amazon Deep Learning AMI GPU TensorFlow 2.7 (Ubuntu 20.04)
- Amazon Deep Learning AMI GPU TensorFlow 2.6 (Amazon Linux 2)
- Amazon Deep Learning AMI GPU TensorFlow 2.6 (Ubuntu 20.04)
- Amazon Deep Learning AMI Graviton GPU TensorFlow 2.6 (Ubuntu 20.04)
- Amazon Deep Learning AMI GPU TensorFlow 2.6 (Ubuntu 18.04)
- Amazon Deep Learning AMI GPU TensorFlow 2.5 (Amazon Linux 2)
- Amazon Deep Learning AMI GPU TensorFlow 2.5 (Ubuntu 20.04)

Multi Framework

GPU

• Amazon Deep Learning AMI (Ubuntu 18.04)

Important NVIDIA driver changes to DLAMIs

On November 15, 2023, Amazon made important changes to Amazon Deep Learning AMIs (DLAMI) related to the NIVIDA driver that DLAMIs use. For information about what changed and whether it affects your usage of DLAMIs, see <u>DLAMI NVIDIA driver change FAQs</u>.

DLAMI NVIDIA driver change FAQs

- What changed?
- Why was this change required?
- Which DLAMIs did this change affect?
- What does this mean for you?
- Is there any loss of functionality with the newer DLAMIs?
- Did this change affect Deep Learning Containers?

What changed?

We split DLAMIs into two separate groups:

- DLAMIs that use NVIDIA proprietary driver (to support P3, P3dn, G3)
- DLAMIs that use NVIDIA OSS driver (to support G4dn, G5, P4, P5)

As a result, we created new DLAMIs for each of the two categories with new names and new AMI IDs. These DLAMIs are *not* interchangeable. That is, DLAMIs from one group don't support instances that the other group supports. For example, the DLAMI that supports P5 doesn't support G3, and the DLAMI that supports G3 doesn't support P5.



Deep Learning Base AMI OSS Nvidia driver

Why was this change required?

Previously, DLAMIs for NVIDIA GPUs included a proprietary kernel driver from NVIDIA. However, the upstream Linux kernel community accepted a change that isolates proprietary kernel drivers, such as the NVIDIA GPU driver, from communicating with other kernel drivers. This change disables GPUDirect RDMA on P4 and P5 series instances, which is the mechanism that allows GPUs to efficiently use EFA for distributed training. As a result, DLAMIs now use OpenRM driver (NVIDIA open source driver), linked against the open source EFA drivers to support G4dn,G5, P4, and P5. However, this OpenRM driver doesn't support older instances (such as P3 and G3). Therefore, to ensure that we continue to provide current, performant, and secure DLAMIs that support both instance types, we split DLAMIs into two groups: one with the OpenRM driver (that supports G4dn, G5, P4, and P5), and one with the older proprietary driver (that supports P3, P3dn, and G3).

Which DLAMIs did this change affect?

This change affected all DLAMIs.

What does this mean for you?

All DLAMIs will continue to provide functionality, performance, and security as long as you run them on a supported Amazon Elastic Compute Cloud (Amazon EC2) instance type. To determine the EC2 instance types that a DLAMI supports, check the release notes for that DLAMI, and then look for **Supported EC2 Instances**. For a list of currently supported DLAMI options and links to their release notes, see <u>Deep Learning AMIs Release Notes</u>.

Additionally, you must use the correct Amazon Command Line Interface (Amazon CLI) commands to invoke the current DLAMIs.

For base DLAMIs that support P3, P3dn, and G3, use this command:

```
aws ec2 describe-images --region us-east-1 --owners amazon \
--filters 'Name=name,Values=Deep Learning Base Proprietary Nvidia Driver AMI (Amazon
Linux 2) Version ??.?' 'Name=state,Values=available' \
--query 'reverse(sort_by(Images, &CreationDate))[:1].ImageId' --output text
```

For base DLAMIs that support G4dn, G5, P4, and P5, use this command:

```
aws ec2 describe-images --region us-east-1 --owners amazon \
--filters 'Name=name,Values=Deep Learning Base OSS Nvidia Driver AMI (Amazon Linux 2)
Version ??.?' 'Name=state,Values=available' \
--query 'reverse(sort_by(Images, &CreationDate))[:1].ImageId' --output text
```

Is there any loss of functionality with the newer DLAMIs?

No, there is no loss of functionality. The current DLAMIs provide all the functionality, performance, and security of the previous DLAMIs, provided you run them on a supported EC2 instance type.

Did this change affect Deep Learning Containers?

No, this change didn't affect Amazon Deep Learning Containers, because they don't include the NVIDIA driver. However, make sure to run Deep Learning Containers on AMIs that are compatible with the underlying instances.

Related information about DLAMI

You can find other resources with related information about DLAMI outside of the Amazon Deep Learning AMIs Developer Guide. On Amazon Web Services re:Post, check out questions about DLAMI from other customers, or ask your own questions. On the Amazon Machine Learning Blog and other Amazon blogs, read official posts about DLAMI.

Amazon Web Services re:Post

Tag: Amazon Deep Learning AMIs

Amazon Blog

- Amazon Machine Learning Blog | Category: Amazon Deep Learning AMIs
- <u>Amazon Machine Learning Blog</u> | Faster training with optimized TensorFlow 1.6 on Amazon EC2 C5 and P3 instances
- <u>Amazon Machine Learning Blog | New Amazon Deep Learning AMIs for Machine Learning</u>
 Practitioners
- <u>Amazon Partner Network (APN) Blog | New Training Courses Available: Introduction to Machine</u> Learning & Deep Learning on Amazon
- Amazon News Blog | Journey into Deep Learning with Amazon

Deprecated features of DLAMI

The following table lists the deprecated features of Amazon Deep Learning AMIs (DLAMI), the date that we deprecated them, and details about why we deprecated them.

Feature	Date	Details
Ubuntu 16.04	10/07/2021	Ubuntu Linux 16.04 LTS reached the end of its five-year LTS window on April 30, 2021 and is no longer supported by its vendor. There are no longer updates to the Deep Learning Base AMI (Ubuntu 16.04) in new releases as of October 2021. Previous releases will continue to be available.
Amazon Linux	10/07/2021	Amazon Linux is <u>end-</u> <u>of-life</u> as of December 2020. There are no longer updates to the Deep Learning AMI (Amazon Linux) in new releases as of October 2021. Previous releases of the Deep Learning AMI (Amazon Linux) will continue to be available.
Chainer	07/01/2020	Chainer has announced <u>the end of major releases</u> as of December, 2019. Consequently, we will no longer include Chainer

Feature	Date	Details
		Conda environments on the DLAMI starting July 2020. Previous releases of the DLAMI that contain these environments will continue to be available . We will provide updates to these environments only if there are security fixes published by the open source community for these frameworks.
Python 3.6	06/15/2020	Due to customer requests, we are moving to Python 3.7 for new TF/MX/PT releases.
Python 2	01/01/2020	The Python open source community has officially ended support for Python 2. The TensorFlow, PyTorch, and MXNet communities have also announced that
		TensorFlow 1.15, TensorFlo w 2.1, PyTorch 1.4, and MXNet 1.6.0 releases will be the last ones supporting Python 2.

Document history for DLAMI

The following table provides a history of recent DLAMI releases and related changes to the Amazon Deep Learning AMIs Developer Guide.

Recent changes

Change	Description	Date
<u>Using TensorFlow Serving to</u> <u>Train a MNIST Model</u>	An example for using Tensorflow serving to train MNIST model.	February 14, 2025
ARM64 DLAMI	The Amazon Deep Learning AMIs now supports images on Arm64 processor-based GPUs.	November 29, 2021
<u>TensorFlow 2</u>	The Deep Learning AMI with Conda now comes with TensorFlow 2 with CUDA 10.	December 3, 2019
<u>Amazon Inferentia</u>	The Deep Learning AMI now supports Amazon Inferenti a hardware and the Amazon Neuron SDK.	December 3, 2019
Installing PyTorch from a Nightly Build	A tutorial was added that covers how you can uninstall PyTorch, then install a nightly build of PyTorch on your Deep Learning AMI with Conda.	September 25, 2018
<u>Conda Tutorial</u>	The example MOTD was updated to reflect a more recent release.	July 23, 2018

Earlier changes

The following table provides a history of earlier DLAMI releases and related changes prior to July 2018.

Change	Description	Date
TensorFlow with Horovod	Added a tutorial for training ImageNet with TensorFlow and Horovod.	June 6, 2018
Upgrading guide	Added the upgrading guide.	May 15, 2018
New regions and new 10 minute tutorial	New regions added: US West (N. California), South America, Canada (Central), EU (London), and EU (Paris). Also, the first release of a 10-minute tutorial titled: "Getting Started with Deep Learning AMI".	April 26, 2018
Chainer tutorial	A tutorial for using Chainer in multi-GPU, single GPU, and CPU modes was added. CUDA integration was upgraded from CUDA 8 to CUDA 9 for several frameworks.	February 28, 2018
Linux AMIs v3.0, plus introduction of MXNet Model Server, TensorFlow Serving, and TensorBoard	Added tutorials for Conda AMIs with new model and visualization serving capabilit ies using MXNet Model Server v0.1.5, TensorFlow Serving v1.4.0, and TensorBoard v0.4.0. AMI and framework CUDA capabilities described in Conda and CUDA overviews. Latest release notes moved to	January 25, 2018

Change	Description	Date
	https://aws.amazon.com/re leasenotes/	
Linux AMIs v2.0	Base, Source, and Conda AMIs updated with NCCL 2.1. Source and Conda AMIs updated with MXNet v1.0, PyTorch 0.3.0, and Keras 2.0.9.	December 11, 2017
Two Windows AMI options added	Windows 2012 R2 and 2016 AMIs released: added to AMI selection guide and added to release notes.	November 30, 2017
Initial documentation release	Detailed description of change with link to topic/sec tion that was changed.	November 15, 2017