

- RHEL9 PyTorch** 3
- 3
- PyTorch** 3
- 6
- 6

RHEL9 PyTorch

RHEL9 PyTorch

PyTorch

NVIDIA Driver, CUDA Driver
[NVIDIA Driver install](#)

가

PyTorch

- <https://pytorch.org/get-started/locally/>

가

START LOCALLY

Select your preferences and run the install command. Stable represents the most currently tested and supported version of PyTorch. This should be suitable for many users. Preview is available if you want the latest, not fully tested and supported, builds that are generated nightly. Please ensure that you have **met the prerequisites below (e.g., numpy)**, depending on your package manager. Anaconda is our recommended package manager since it installs all dependencies. You can also [install previous versions of PyTorch](#). Note that LibTorch is only available for C++.

PyTorch Build	Stable (1.13.1)		Preview (Nightly)	
Your OS	Linux	Mac	Windows	
Package	Conda	Pip	LibTorch	Source
Language	Python		C++ / Java	
Compute Platform	CUDA 11.6	CUDA 11.7	ROCm 5.2	CPU
Run this Command:	<pre>pip3 install torch torchvision torchaudio --extra-index-url https://download.pytorch.org/whl/cu116</pre>			

NOTE: PyTorch LTS has been deprecated. For more information, see [this blog](#).

OS,

, CUDA
가

Run this Command

```
[root@kvm33 ~]# pip3 install torch torchvision torchaudio --extra-index-url
```

```
https://download.pytorch.org/whl/cu116
Looking in indexes: https://pypi.org/simple,
https://download.pytorch.org/whl/cu116
Collecting torch
  Downloading
https://download.pytorch.org/whl/cu116/torch-1.13.1%2Bcu116-cp39-cp39-linux_
x86_64.whl (1977.9 MB)
----- 2.0/2.0 GB 1.5 MB/s eta
0:00:00
Collecting torchvision
  Downloading
https://download.pytorch.org/whl/cu116/torchvision-0.14.1%2Bcu116-cp39-cp39-
linux_x86_64.whl (24.2 MB)
----- 24.2/24.2 MB 10.6 MB/s eta
0:00:00
Collecting torchaudio
  Downloading
https://download.pytorch.org/whl/cu116/torchaudio-0.13.1%2Bcu116-cp39-cp39-l
inux_x86_64.whl (4.2 MB)
----- 4.2/4.2 MB 10.2 MB/s eta
0:00:00
Collecting typing-extensions
  Downloading typing_extensions-4.4.0-py3-none-any.whl (26 kB)
Collecting pillow!=8.3.*,>=5.3.0
  Downloading Pillow-9.4.0-cp39-cp39-manylinux_2_28_x86_64.whl (3.4 MB)
----- 3.4/3.4 MB 9.9 MB/s eta
0:00:00
Collecting numpy
  Downloading numpy-1.24.1-cp39-cp39-
manylinux_2_17_x86_64.manylinux2014_x86_64.whl (17.3 MB)
----- 17.3/17.3 MB 10.8 MB/s eta
0:00:00
Requirement already satisfied: requests in /usr/lib/python3.9/site-packages
(from torchvision) (2.25.1)
Requirement already satisfied: chardet<5,>=3.0.2 in /usr/lib/python3.9/site-
packages (from requests->torchvision) (4.0.0)
Requirement already satisfied: idna<3,>=2.5 in /usr/lib/python3.9/site-
packages (from requests->torchvision) (2.10)
Requirement already satisfied: urllib3<1.27,>=1.21.1 in
/usr/lib/python3.9/site-packages (from requests->torchvision) (1.26.5)
Installing collected packages: typing-extensions, pillow, numpy, torch,
torchvision, torchaudio
Successfully installed numpy-1.24.1 pillow-9.4.0 torch-1.13.1+cu116
torchaudio-0.13.1+cu116 torchvision-0.14.1+cu116 typing-extensions-4.4.0
WARNING: Running pip as the 'root' user can result in broken permissions and
conflicting behaviour with the system package manager. It is recommended to
use a virtual environment instead: https://pip.pypa.io/warnings/venv
[root@kvm33 ~]# python
Python 3.9.14 (main, Nov 7 2022, 00:00:00)
```

```
[GCC 11.3.1 20220421 (Red Hat 11.3.1-2)] on linux
Type "help", "copyright", "credits" or "license" for more information.
>>> import torch
>>> [Ctrl+D]
```

```
import torch
import math

print(torch.__version__) # torch version

dtype = torch.float
# device = torch.device("cpu")
device = torch.device("cuda") # Uncomment this to run on GPU, GPU

# Create random input and output data
x = torch.linspace(-math.pi, math.pi, 2000, device=device, dtype=dtype)
y = torch.sin(x)

# Randomly initialize weights
a = torch.randn((), device=device, dtype=dtype)
b = torch.randn((), device=device, dtype=dtype)
c = torch.randn((), device=device, dtype=dtype)
d = torch.randn((), device=device, dtype=dtype)

learning_rate = 1e-6
for t in range(2000):
    # Forward pass: compute predicted y
    y_pred = a + b * x + c * x ** 2 + d * x ** 3

    # Compute and print loss
    loss = (y_pred - y).pow(2).sum().item()
    if t % 100 == 99:
        print(t, loss)

    # Backprop to compute gradients of a, b, c, d with respect to loss
    grad_y_pred = 2.0 * (y_pred - y)
    grad_a = grad_y_pred.sum()
    grad_b = (grad_y_pred * x).sum()
    grad_c = (grad_y_pred * x ** 2).sum()
    grad_d = (grad_y_pred * x ** 3).sum()

    # Update weights using gradient descent
    a -= learning_rate * grad_a
    b -= learning_rate * grad_b
    c -= learning_rate * grad_c
    d -= learning_rate * grad_d
```

```
print(f'Result: y = {a.item()} + {b.item()} x + {c.item()} x^2 + {d.item()} x^3')
```

test-torch.py

test-torch.py

torch.py

torch.py

가

import torch

```
[root@kvm33 ~]# python
Python 3.9.14 (main, Nov 7 2022, 00:00:00)
[GCC 11.3.1 20220421 (Red Hat 11.3.1-2)] on linux
Type "help", "copyright", "credits" or "license" for more information.
>>> import torch
Traceback (most recent call last):
  File "<stdin>", line 1, in <module>
  File "/root/torch.py", line 4, in <module>      ##### <<<<<<
    /root/torch.py
    print(torch.__version__) # torch version
AttributeError: partially initialized module 'torch' has no attribute
'__version__' (most likely due to a circular import)
>>>
```

- <https://pytorch.org/get-started/locally/>
- <https://acton21.tistory.com/13>

From: <https://at1.kr/dokuwiki/> - AllThatLinux!

Permanent link: https://at1.kr/dokuwiki/doku.php/rhel9_pytorch_%ED%99%98%EA%B2%BD_%EA%B5%AC%EC%84%B1?rev=1672898052

Last update: 2023/01/05 05:54

