

Ubuntu 24.04 Upgrade

Kang Lu

2024 - 10 - 09

Today I decided to upgrade one of my servers to Ubuntu 24.04.01. The current server was running Ubuntu 22.04.5 LTS. This server is primarily used for video encoding. The server has an NVIDIA GTX 1060 GPU installed, using the GPU for HEVC video encoding is about 10 times faster than on a regular CPU.

I thought all I had to do was:

```
sudo apt update
sudo apt dist-upgrade
sudo reboot
sudo do-release-upgrade
```

1
2
3
4

Listing 1: shell commands to upgrade

There were many packages that needed to be removed, installed. This took nearly an hour. There were a couple of configuration files concerning firmware update scheduling and grub

that require manual interventions. I decided to adopt the new versions and forego the previous versions, since I did not make any previous modifications.

Once the upgrade was completed, I had to reboot the machine. All the basic things continue to work. The GUI Desktop Environment booted up without any issues, and most of the things were there and operational. However, there was one thing that was missing.

When I execute `nvidia-smi`, it indicated to me that no GPU was detected. This means the original CUDA Toolkit is no longer recognized.

I had to follow the instructions in Listing 2 which was detailed in <https://developer.nvidia.com>.

```
wget https://developer.download.nvidia.com/compute/cuda/repos/ubuntu2404/x86_64/cuda-ubuntu2404.pin
sudo mv cuda-ubuntu2404.pin /etc/apt/preferences.d/cuda-repository-pin-600
wget https://developer.download.nvidia.com/compute/cuda/12.6.2/local_installers/cuda-repo-ubuntu2404-12-6-local_12.6.2-560.35.03-1_amd64.deb
sudo dpkg -i cuda-repo-ubuntu2404-12-6-local_12.6.2-560.35.03-1_amd64.deb
sudo cp /var/cuda-repo-ubuntu2404-12-6-local/cuda-*-keyring.gpg /usr/share/keyrings/
sudo apt-get update
sudo apt-get -y install cuda-toolkit-12-6
sudo apt-get install -y cuda-drivers
```

1
2
3
4
5
6
7
8

Listing 2: commands to install CUDA Toolkit

After a quick reboot, the GPU is recognized again with the CUDA Toolkit. However when I tried to run `ffmpeg` I got a missing share library error. This makes sense, since the previous shared library may have been replaced with the upgrade, and I did not statically build `ffmpeg`.

```
./configure --enable-nonfree --enable-gpl --enable-libzimg \  
--enable-libx265 --enable-libx264 --enable-cuda-nvcc \  
--enable-libnpp \  
--extra-cflags=-I/usr/local/cuda/include \  
--extra-ldflags=-L/usr/local/cuda/lib64 \  
--disable-static --enable-shared --enable-libass
```

1
2
3
4
5
6

Listing 3: custom configure with more options

I had to follow the compilation instructions on <https://docs.nvidia.com>. There was one exception that I had to make. During the `configure` step, I had to add a few more options so that `x264` codec is also available. See Listing 3.

I also ran into an issue where `nasm` was not installed so I had to install that before the `configure` step.

Overall it took a little over an hour to complete the upgrade process.