

2021

- January
- LBA write rate on database host
- March
- April
- December

January

1/15/2021

Upgraded the following machines from debian 9 to debian 10:

- postgres
- brat-anno

1/18/2021

When removing the head unit in my car, the head unit is attached to a carriage with screws. Those are the screws you see. When removing the head unit though, you don't have to worry about these screws: you just need to slide the carriage out. There are little tabs that need to be pushed in so that the whole thing can slide out. The top tabs are probably already pushed in, they're quite worn at this point.

LBA write rate on database host

Date	LBAs written
3/6/2028	8576402139
3/13/2028	679613123
3/31/2029	967961811

.18 TB over 25 days. 160 TBW for the drive.

Should last another 2.5 years.

March

3/18/2021: The cooler on the 7900x needs to be shifted over so that it covers the IHS more completely before you run linpack etc. on it. Ram slots A1, A2, B1, B2 are the ones that work (e.g. all ram slots left of cpu socket).

April

An automatic update to cuda has broken my virtual environments. Creating new virtual environments doesn't seem to help. I get an error saying that all available cuda devices are busy when there are actually no devices in use. (e.g. nvidia-smi shows 0% gpu utilization, no vram occupied and no processes claiming the gpu).

In addition, snapper has non-sane defaults on ubuntu and thus my rollback was broken. After rolling back, the .snapshots folder was lost so I was unable to rollback again. I found a solution which involves writing something to your fstab for .snapshots. I copied the actual config from my opensuse tumbleweed install and it worked. However, no matter how far back I go with my snapshots, I still get the cuda device busy error. :(.

I booted another ubuntu os from my portable ssd with cuda installed on it and sure enough it works, I can train the model (and see both the gpus in nvidia-smi).

I am going to try and see if my m40 gpu is compatible with the cuda drivers on opensuse leap (Hopefully this is the case as I've purchased another m40 for cheap to put in my opensuse leap system). If, indeed it does work, I will be switching to opensuse leap for mwanafunzi. I know that snapper has good defaults and works well on tumbleweed and leap so hopefully this rollback issue won't be a problem in the future and if nvidia breaks my setup, I can just rewind. This was the original reason why I was using snapper but apparently I had it misconfigured.

December

12/10/2021

- Since we moved, everything had to be migrated to comcast's internet.
- All machines were down from 12/7 to 12/9
- Had to change ip for internal network on router from 10.0.0.X/32 to 192.168.X.X to match previous ip range
 - IP addresses for cluster nodes cannot change so we have to keep the internal network ips the same.
- I'm now using a tiny linode instance, using wireguard linked back the nginx reverse proxy
 - Largely followed [this guide](#), which requires some initial steps from [this guide](#).
 - Private and public wireguard keys generated using [this guide](#).
 - Seems to work well but jellyfin can't use large files since my upload speed on comcast is horrible.
- Modified google domains config to remove the dynamic dns entry and instead switched to a static A record.