

Kubernetes Cluster

Lester Vecsey

1 Background

The basis of this project is a Ceph FS (filesystem) and a reasonable high speed network.

The actual project is a master server with a dual port 100 GbE network card, and two worker nodes that are also capable of 100 GbE network connectivity.

A 10 GbE network card serves as an uplink to the rest of the network.

2 Cluster specs

| | Cores | Memory (Gi) |
|-----------|-------|-------------|
| Master | 24 | 32 |
| Worker #1 | 24 | 32 |
| Worker #2 | 24 | 32 |
| Totals | 72 | 96 |

The motherboard supports more memory. The height of the RAM modules needs to be lower in order to have four DIMMs installed, and to clear the CPU Cooler.

3 Master computer

| Item | Model | Notes or Comments |
|--------------------------------------|---|--------------------------|
| Chassis | Fractal Design Define 7 Mini | NewEggBusiness - 15.98" |
| Motherboard | Asus TUF GAMING B760M-PLUS II | NewEggBusiness |
| Memory | Corsair DDR5 6600 | NewEggBusiness |
| M.2 Storage | Team Group MP44 M.2 2280 1TB PCIe | NewEggBusiness |
| Boot Storage #1 | Crucial MX500 (500 GB) | NewEggBusiness |
| Boot Storage #2 | Crucial MX500 (500 GB) | NewEggBusiness |
| 10 GbE NIC | TP-Link TX401 | NewEggBusiness |
| CPU | Intel Core i9-14900K (or the KS) w/ integrated video | NewEggBusiness |
| PSU | Corsair RM650 | Amazon |
| CPU Cooler | Noctua NH-U12a chromax.black | Amazon |
| 10ft Red CAT6 cable | | Amazon - used for uplink |
| 100 GbE NIC | Intel E810-CQDA2 | ServerOrbit |
| Two DAC cables (2m length preferred) | | Naddod |

You can also use fiber optic cables in place of the DAC cables.

I could have done this if I were to place the master node closer to a top shelf, and then the worker nodes could remain on the bottom shelf.

An uninterruptable power supply by CyberPower, model **CP1500PFCLCD** was also added to the list. It will power all three servers.

Also required was a 5-tier shelf from Temu, to hold everything. A Wi-Fi printer was optional.

4 Worker computers

Similar to the above list, however the only storage drive is **M.2 SAMSUNG 990 EVO SSD 2TB, PCIe Gen 4x4**.

The CPU is Intel Core **i9-14900KF** without integrated video.

There are no dedicated boot drives because the M.2 will have a partition for that.

A **single port Intel 100 GbE** network card would be sufficient however I went with the same card as in the master server, so a dual-port.

There is no TP-Link TX401 in this one, since the only uplink needed is the 100 GbE link.

For the chassis I went with more of a cube shape, with the **Node 804** model.

In addition, each worker node needed a video card. It has a **PCI x4** interface which uses an appropriate slot on the motherboard.

5 Additional notes

The CPU Cooler comes with paste for applying to the chip.

It was somewhat time consuming to assemble everything, however for the most part it went smoothly.

The parts had to be reworked a bit because at first I wasn't sure that I only had CPU's without integrated video.

If you can go with all integrated video, that would probably be the better choice.

6 Computer Assembly

- Assemble the Motherboard, Memory, M.2 Storage, CPU, and CPU Cooler as one complete module.
- Attached the needed cables to the the Power Supply Unit (PSU), installing it into the case with 4 screws.
- Screw in and tighten the needed motherboard standoffs into the Micro ATX portion of the case.
- Insert the motherboard assembly into the case, and tighten down with M3 screws.
- Attach main cables, CPU fans, and front connectors such as power led, hdd led, and power switch.
- Install any networking card(s) or video card (optional)

7 Network wiring

DAC cables are used for interconnects between the master server, and each worker node. Thus no intermediary switch is actually needed, for this very tiny or minimal cluster configuration.

8 Software configuration

Operating System is Ubuntu 24.04.1 which is installed from a small, 4GB sized USB thumb drive.

Using an 8GB size thumb drive or larger will be easier to flash from more places.

9 Build Log

| Date | Computer | Notes |
|-----------|-----------|---|
| 9/25/2024 | Worker #1 | Functioning, just needs high speed networking. |
| 9/25/2024 | Worker #2 | Functioning, just needs high speed networking. |
| 9/27/2024 | Master | Assembly and OS installation w/ on-board video |
| 10/2/2024 | All | Install remaining network cards and connect with cables |

10 Initial configuration

10.1 Networking

The computers each need to reach each other. I used the following hostnames:

| Hostname | IP Address |
|----------------------------|---------------|
| kubmaster-0.local | 192.168.2.70 |
| kubworker-internal-0.local | 192.168.10.73 |
| kubworker-internal-1.local | 192.168.11.74 |

10.1.1 Netplan

Netplan configuration files are in `/etc/netplan` for each system.

10.2 NTP

Setting up time synchronization for all three systems.

The configuration file is located at `/etc/systemd/timesyncd.conf`

For kubmaster-0.local

```
[Time]
NTP=mds2.local
```

For the first worker:

```
[Time]
NTP=192.168.10.70
```

For the second worker:

```
[Time]
NTP=192.168.11.70
```

10.3 Kubernetes

In the next article the software will be installed from <https://kubernetes.io/>

11 Additional

11.1 Container registry

An additional small form factor computer, a Raspberry PI 5 can be added next to the master node. A 10ft LAN cable connects to the same switch as the uplink cable from the master node. A 512GB SD card such as Sandisk Extreme can be used, with a transfer rate of 190 MB/s. Alternatively, if just getting started then you can use one of the Raspberry PI 5 kits that include an SD card such as 128 GB.

It runs a local instance of the <https://quay.io/> software.

- Operating System: Ubuntu 24.04.1 LTS (Noble Numbat)
- Hostname: registry.local
- IP Address: 192.168.2.68

11.1.1 Packages

```
sudo apt-get install openssh-server nut-client
```

12 Maintenance

13 Future directions