

SubMIT

Recent & Future Upgrades

Zhangqier Wang
SubMIT Review
Jul. 18th 2023







Outline



- Recent Upgrades
 - Network
 - IPv6 address
 - Filer (/home & /work)
 - New machines
 - Fast mount space
- Future: Linux Upgrade
 - Plans
 - Mitigate user disruption

- 100 Gb/s link is deployed on all the major machines
 - 9 servers: the login pool for SubMIT
 - 10 servers: user storage system (gluster filesystem)
 - machines with special services: fast connection mount (scratch), powerful machines, other services
- Main network switch 100 Gb/s link
 - 100 Gb/s network between Building 24 and Bates
- A high-speed link that is capable of transmitting large amounts of data in a short period within SubMIT systems.
 - In CMS analysis context: the first 100 Gb/s linked analysis facility



Yellow and pink cables are 100 Gb/s links

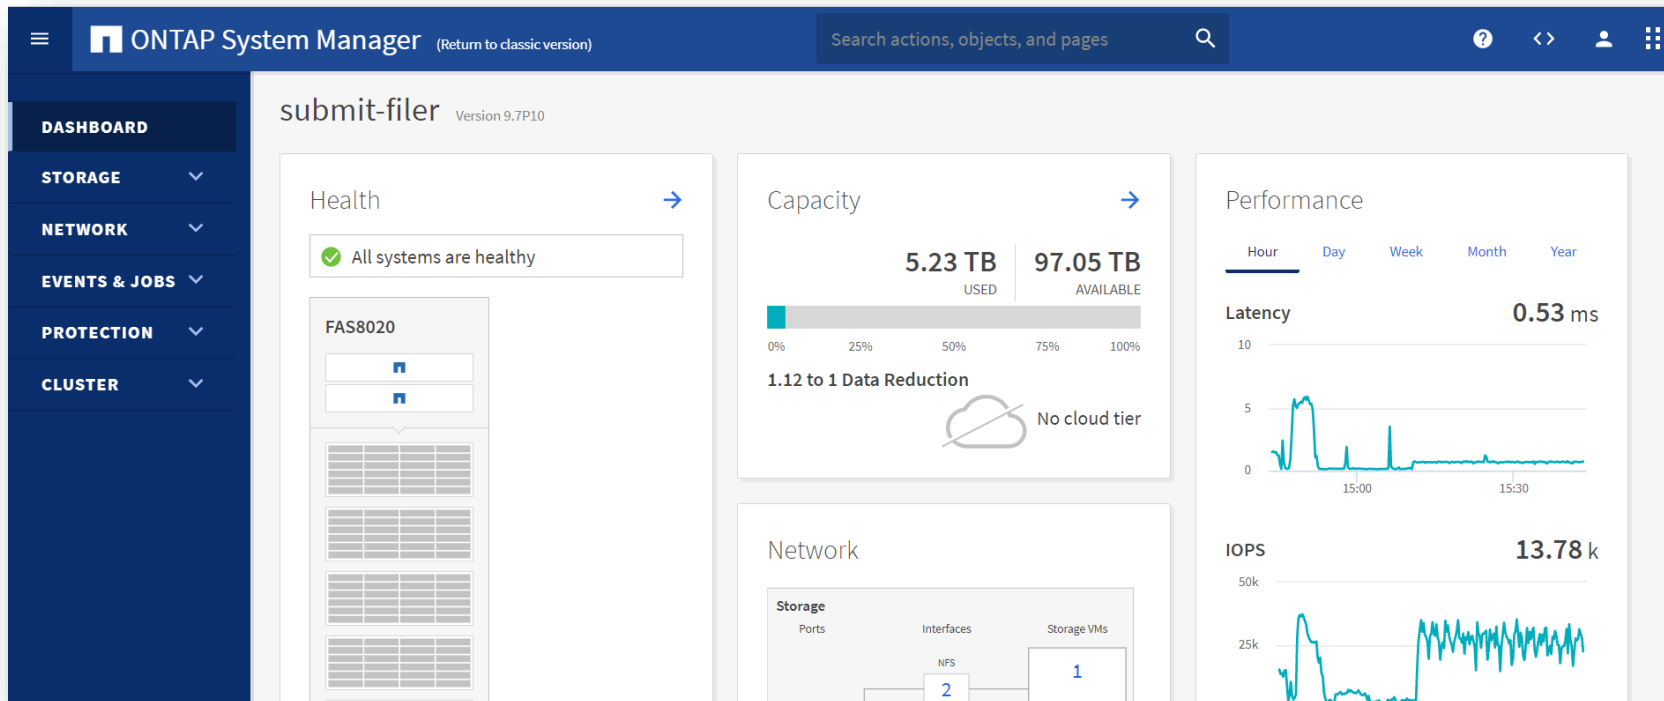
- IPv6 is the next generation Internet Protocol (IP) standard intended to eventually replace IPv4
- Dual IP connections are enabled in the SubMIT system
 - Each machine has supported both IPv6 and IPv4 network.
 - Dual IP address added to the DNS
 - The IPv6 routing is correctly set up






Filer

- The filer is used to manage the user space
 - dedicated storage devices provide file-level data storage to multiple clients or users
- In total ~100 TB space for user /home and /work



- Two powerful interactive machines added to SubMIT
 - 192 cores/384 threads, memory 4 GB/thread
- Ten machines for storage
 - Distributed storage system (Gluster filesystem)
 - spinning disks
 - better performance in data storage
 - 1 TB each user by default
 - Support remote connection
 - Total space: 392 TB
 - Part of Slurm pool: total 640 threads
- GPU machines
 - Refurbished LQCD machines
 - 10 machines with 40 GPUs ([GeForce GTX 1080 Graphics Cards](#))



Storage system servers



Power Upgrades



- Most SubMIT machines are located in the computing room 24-032 (Building 24)
- The power circuit was upgraded to provide 208 V to the SubMIT servers.
 - Enough power to support the new machines
- Building 24 has stable power from campus

News, Project Team

Power circuit work completed

by Zhanqier Wang - November 16, 2022

On November 16, 2022 from 5:00-8:00am, during the scheduled subMIT downtime, the new power circuit carrying the required 208 V was [...]

[Continue reading »](#)





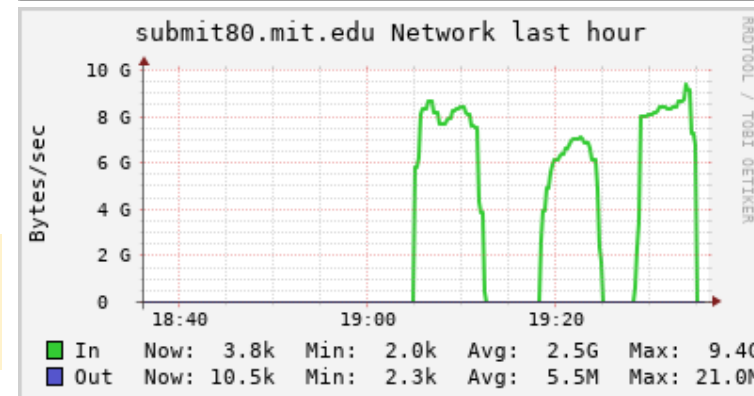
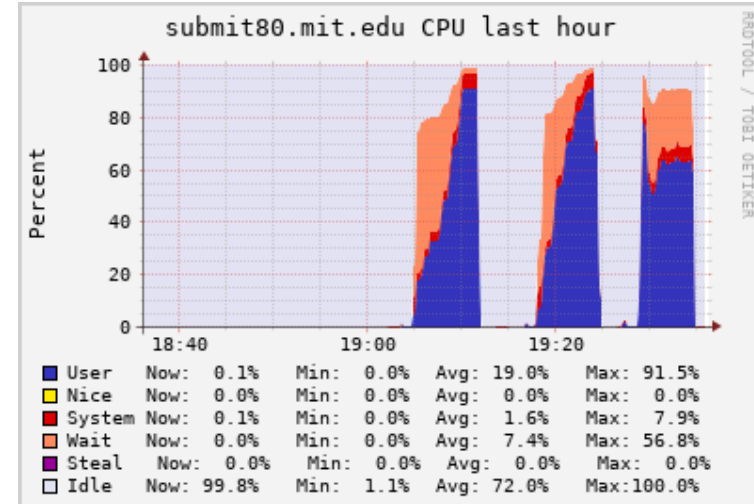
Scratch Mount



- A fast mount space mounted as /scratch
- Hosted by a dedicated nfs server
 - four NVMe Gen4 disks
 - 28 TB space
 - combined in md RAID0
 - ext4 file system
 - shared over Network File System
 - 64 threads to handle incoming NFS requests
- On client side:
 - nfsrahead nfs4=2048
 - Enhance the read performance

Real life data analysis using samples from scratch with one of the new high core machines (submit80)

Reach a maximum of 75Gbits/sec (9.4 GB/s)!





Operating System

- AlmaLinux 9 will be the next generation operating system, replacing the current Centos 7 system
 - Long-term support for AlmaLinux 9
 - Improved performance and features, better user experience
 - Enhanced security
 - Continuity and compatibility, easy to maintain as it is downstream of RHEL.
- The upgrade is conducted with a testing machine.
- A few submit machines are upgraded/installed in alma9 with migrated services
 - Migrating condor service
 - Slurm AlmaLinux 9 pool
 - [Documentation](#)



Future Upgrade

- Upgrade all computers operating systems to AlmaLinux9
- Plans
 - Upgrade the system in stages
 - Testing IP/hostname reserved in DNS
 - Slurm partition for Alma 9 open to user
 - Adjustment based on user feed back
- Disruption mitigation
 - Containers of Centos 7 are available in the cvmfs area
 - Dedicated containers are available to support LHC experiments research
 - old environments are accessible through singularity/conda environments

[Documentation](#)


Move to AlmaLinux

In December 2020, Red Hat announced that development of CentOS, a free-of-cost downstream fork of the commercial Red Hat Enterprise Linux (RHEL), would be discontinued. As such, we will move to AlmaLinux, a community-supported, production-grade enterprise operating system that is binary-compatible with RHEL. This page covers the planned upgrade and testing of AlmaLinux on submit.

How we plan to upgrade submit

Because it is hard to predict how an operating system upgrade will affect everyone's work, we will upgrade the system in stages and encourage users to test their workflows on AlmaLinux. For now, the submit system uses CentOS 7 with AlmaLinux machines available only through slurm using a test partition. For the future upgrades, the submit machines will get upgraded with announcements in the submit-users email group.

- Several upgrades have been conducted in the past few months
 - Network, hardware, computing resources, operating system
 - Achieve a more powerful and reliable analysis facility
- Operating system will be upgraded to AlmaLinux 9
 - Plans are made to ensure a smooth transition and good user experience.
- There are a number of smaller software installation and upgrade projects on going: i.e. Nextcloud, Visual Studio.



The End