

# SubMIT

## Recent & Future Upgrades

Zhangqier Wang  
SubMIT Review  
June. 21<sup>st</sup> 2024





# Outline



- Recent Upgrades
  - Operation System Upgrade
    - Centos 7 -> Alma Linux 9
- Future: Storage system upgrade
  - Mitigation from GlusterFS -> CephFS
  - Current status and plans

- AlmaLinux 9 is the next generation operating system, replacing the current Centos 7 system
  - Centos 7 end-of-life on Jun 30<sup>th</sup> 2024
    - no more security updates or support
  - Compatibility with RHEL 9, easy to maintain, long-term support
  - Enhanced security and performance
    - Updated security features and patches
    - Improved performance with updated kernels and software packages
  - Modern software compatibility





# AlmaLinux 9 Migration

- The Alma 9 Migration includes:
  - 8 login machines
  - 2 powerful interactive machines
    - 192 cores / 384 threads, memory 4 GB / thread
  - Fast mount server
    - Holds a fast mount space
  - Storage system: 10 machines
  - GPU machines
    - 4 [NVIDIA A800](#) GPU machines
    - 10 machines with 40 GPUs ([GeForce GTX 1080 Graphics Cards](#))
- All the machines are upgraded to AlmaLinux 9 and the major services are running fine



# Migration Procedure

- The goal is to have a smooth migration, which causes minimal interference to user experience
  - Upgrade conducted in three stages: preliminary -> partial -> full
- Preliminary stage
  - Alma 9 installation on test machines
  - Finalize the post-installation to ensure it contains the fundamental functions and software
    - Alma 9 Network configuration
    - User information from LDAP
    - /home, /work, storage system mounted
    - Slurm, condor connected to the central server

Preliminary Stage

Partial Migration

Full Migration



# Migration Procedure

- Partial migration
  - Upgrade partial machines without crucial services
  - Deploy crucial services in the testbed and make sure they work
    - Submit website
    - Condor/slurm central services
    - Monitoring system (ganglia)
    - Login load balance
  - For the local computing resources slurm pool, an Alma 9 partition is created
    - We hear the feedback from voluntary users who use Alma 9 partition to make sure the upgrade performs well
- Full migration
  - Once the test is complete, formally migrate the crucial services to the Alma 9 machines
  - Upgrade the rest of the machines



# Centos 7 Singularity



- Users conducted their research under centos 7, their setup may not be compatible with Alma 9
- To avoid causing major delays to users' research progress, we offer a "simple, straightforward" way for users to access the centos 7 environment: centos 7 singularity
- A recommended command line is provided in the user guide to utilize the centos 7 singularity from CVMFS

## Will CentOS still be available

If your workflow is constrained to CentOS, submit will provide a centrally available singularity of CentOS7.9. You can test this singularity below.

```
singularity shell /cvmfs/cvmfs.cmsaf.mit.edu/submit/work/submit/submit-software/centos/centos7
```

Here is the recommended command to use centos 7 singularity

```
singularity shell --bind /cvmfs:/cvmfs --bind /work/submit:/work/submit --bind /data/submit:/d
```

You can add any additional binding (with whatever you are using). Then use "bash" command to load your personal bash setup.

If you need something specific besides this singularity, please email us at [submit-help@mit.edu](mailto:submit-help@mit.edu)

# Storage System Upgrade



- Current storage system is “GlusterFS”, managed by 10 machines
  - Total 40 disks, with total space 720 TB
  - Xrootd system installed on each machine for remote access
- We decided to migrate to storage system “CephFS”
  - GlusterFS is also reaching the end-of-life at the end of 2024
  - Performance limitation during the last year operation
    - when it was overloaded by a certain user, it became extremely slow for others
    - The folder in Gluster may get dismounted when too many files are transferred.




- CephFS suits the SubMIT project better than GlusterFS
  - **Ceph** distributes data across computers in the cluster and allows the user to access all of the data at once through the interface.
    - stores data without much user intervention.
    - Multiple clients can also access the store without intervention.
  - **Gluster** data storage happens in blocks.
    - Performance can degrade with large numbers of small files and it can become less efficient under heavy workloads.
- A test for the T2 storage and it is to improve the data storage performance





# Upgrade Status and Plan

- 10 machines that manage the storage system have been upgrade to Alma Linux 9
- Gluster system and xrootd service are upgraded accordingly and function well
- Hardwares were bought, and extra disks are put in those machines to set up the CephFS storage system besides gluster
  - 12 disks spread out in 10 machines
  - More disks on the way
- Performance test on CephFS, finalizing parameters
- Future Plan
  - Data transferring from Gluster
  - Decommission Gluster disks and integrate them to CephFS






# Summary



- The Alma 9 upgrades have been conducted in the past few months and completed
  - Network, services, softwares, local computing resources
  - Achieve a more powerful and reliable analysis facility
- Storage system will be migrated to CephFS
  - Machines holding the storage system are upgraded to Alam 9
  - CephFS system is installed besides glusterFS for the test
  - Plans are made to ensure a smooth transition and good user experience.



The End