


EOS SE Status at IHEP

Xuantong Zhang,

On Behave of IHEP-CC Storage/Grid Team,

EOS Workshop, 2023



Outline



1. Introduction
2. Storage elements construction and deployment for JUNO Experiment
3. DPM -> EOS upgrading for WLCG Beijing T2 Site

Current SE for JUNO



Introduction



EOS working as Grid storage system at IHEP,

- Since 2021, for JUNO experiment.
- Since April 2023, EOS-CTA for JUNO.
- Since May 2023, for Beijing Grid site storage (Atlas/CMS/BelleII).
- Since May 2023, for new LHCb Beijing Tier1 storage.

April 2021

Disk Storage
as a JUNO
Grid site.

April 2023

Tape Storage
as a JUNO
Grid site.

May 2023

DPM upgrade
to EOS for
Beijing T2
site.

May 2023

A new LHCb
T1 site.

EOS SE for JUNO



Work as a DIRAC and Rucio site of JUNO.


- **Hardware,**
 - 8 nodes, 4 job disk arrays.
- **Software,**
 - EOS server, v4.8.62
 - 1 MGM, 14 FTS nodes, 504 FS in total, 3 QuarkDB nodes.


More details about EOS at IHEP, refer to Haibo' s talk.


Only 1 instance for both local and Grid.




- Local and Grid storage are separated by different paths.
- Grid storage path basically managed by DIRAC/Rucio, directly access by Gfal or other Grid tools are not recommended to users.

EOS Status at IHEP

 2023年4月25日 13:45

 15m

 31/3-004 - IT Amphitheatre (CERN)


  

20 Minutes

Sites and Deploye...

EOS Operation & Sites

报告人

 LI Haibo lihaibo

EOS CTA for JUNO



Still in test but nearly in production.

- **Hardware**
 - 2 nodes, 3 LTO 9 drives
 - 120 LTO 9 tapes
- **Software**
 - CTA v5.7.17
 - EOS v5.1.9

More info. about JUNO CTA, refer to Yujiang's talk.

EOS buffer is an independent instance from disk storage.

- Only serves the JUNO production users from Grid.
- A path is set to automatically archive data for JUNO Grid users.
- Managed as a DIRAC/Rucio tape site.

The screenshot shows a presentation slide with a dark background and orange text. The title is 'CTA Status and Future at IHEP'. Below the title, there is a date and time '2023年4月26日 09:55', a duration '20m', and a location '31/3-004 - IT Amphitheatre (CERN)'. On the right side, there are three buttons: '10 Minutes', 'CTA', and 'CTA'. Below the location, there is a section titled '报告人' (Reporter) with a person icon and the text '博士 Yujiang Bi (Institute of High ...)'.

User Authentication Configs



For GSI authentication,

- User DN mapping list synchronized from VOMS.
- Grid users in same roles are mapped to one same grid user.
- Unsafe, plan to be solved.

```
sec.protocol gsi -crl:0 -cert:/etc/grid-security/daemon/hostcert.pem -key:/etc/grid-security/daemon/hostkey.pem -gridmap:/etc/grid-security/dn-grid-mapfile -moninfo:1 -d:3 -gmapopt:2 -vomsat:1 -moninfo:1 -exppxy:=creds
```

For Token-based authentication,

- Macaroon tokens supported in production server,
- Sci-tokens supported in testbed and CTA buffer and finished test.
- At present, Grid users are also mapped to a single grid user.
- Plan to follow the IAM user list to do fine-grained authentication.

```
mgmofs.macaroonslib /opt/eos/xrootd/lib64/libXrdMacaroons.so /usr/lib64/libEosAccSciTokens.so
```

```
[Global]
audience = https://wlcg.cern.ch/jwt/v1/any

[Issuer JUNO-Connect]
issuer = https://iam-juno.cloud.cnaf.infn.it/
base_path = /
map_subject = False
default_user = sjuno001
```

SE Operation Issues



MGM crash issue,

- **Description:** MGM crashed if no token library specified but a simple HTTP POST request is sent.
- **Status:** Problem reported and got solved, with upgraded to v4.8.60.

MGM crash issue,

- **Description:** MGM crashed if removing a file with invalid characters, which caused null path when assigning to the vid.path.
- **Status:** Problem reported and got solved, with upgraded to v4.8.63.

StoRM to EOS HTTP push mode TPC failure issue,

- **Description:** StoRM to EOS push TPC failure due to the header of HTTP request exceeding 1000 Bytes.
- **Status:** Problem reported and got checked by EOS team, but still in resolving, need to be fixed by StoRM team.

DPM->EOS Upgrades for WLCG experiments



DPM -> EOS Upgrading



Motivation,

- DPM is retiring.
- We have more experience on EOS operation.
- New LHCb Beijing T1 site also chooses EOS.

Resources for new EOS at Beijing site,

- **ATLAS: 400 TB**
 - Poolnodes: 3 pool nodes, DELL R620 10/2014
 - Disk Array: 5 X TONGYOU 24Slots Array
- **CMS: 680 TB**
 - Poolnodes: 2 pool nodes, DELL R640 07/2014
 - Disk Array: DELL ME4084, 12TB*84
- **LHCb: 250 TB**
 - Poolnodes: 1 pool nodes, DELL R630 11/2018
 - Disk Array: DELL MD 3860F
- **BELLEII: 280 TB**
 - Poolnodes: 1 pool nodes DELL R740 07/2019,
 - Disk Array: DELL ME4084



VO Data Strategy

No new storage, so old machines are reused.

Central managed data get drained by experiment VOs, user data get moved by users themselves.

ATLAS

- DATADISK 60TB will be moved by ATLAS DDM.
- LOCALGROUPDISK 27TB stored by user.

CMS

- Data moved by CMS VO.
- User data store by user.

LHCb

- DATA moved to IHEP Tier1.

Belle II

- Notification to VO

EOS SE Configuration



1 instance for all VOs.

- Xrootd with GSI authentication.
- HTTP with macaroon and scitokens.

**ccsrm.ihep.ac.cn ->
cceos.ihep.ac.cn**

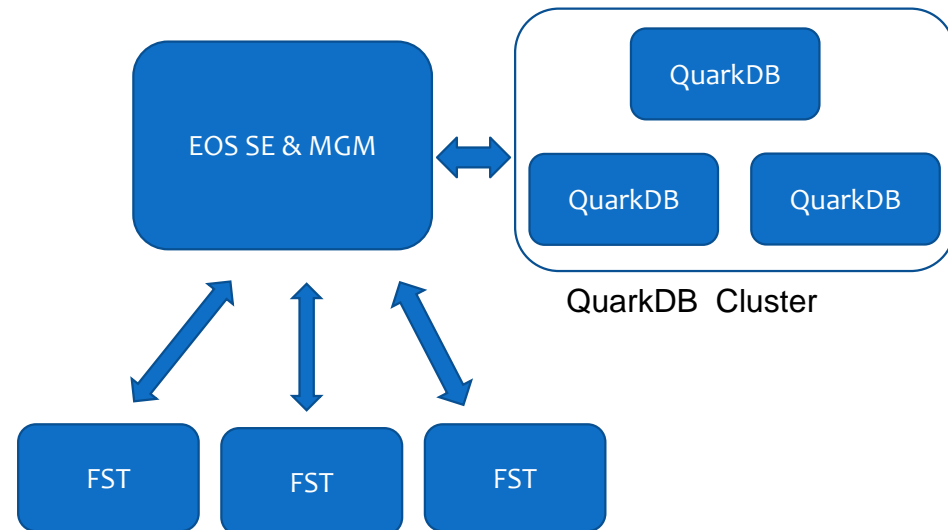
- /eos/ihep/atlas
- /eos/ihep/cms
- /eos/ihep/belle
- LHCb data to new T1 storage.

For each VO,

- Independent space and FTS.

Timeline,

- Moving data, since March 2023,
- Downtime for upgrading, 8th May to 10th May.



Conclusion and Plan



EOS works as Grid storage at IHEP for 2 years.

- Works as DIRAC/Rucio site for Grid users at JUNO experiment.
- Ready with macaroon tokens and start to move to Sci-tokens.
- Some naïve mapping needs to be get fine-grained in future.

EOS will replace retiring DPM as future T1/T2 storage of WLCG at IHEP this year.

- Reuse all old machines from DPM to EOS.
- 1 instance endpoint but with different nodes for different VOs.
- VOMS and tokens plan to be supported but need more careful configuration.

Thank You!

