

12:30 pm Central

URL: <https://unl.zoom.us/j/183382852>

Phone: +1 669 900 6833 or +1 408 638 0968 or +1 646 876 9923

Meeting ID: 183 382 852 (password required)

Operations Meeting October 10, 2025

Attendees: Jeff D, Ashton, John, Huijun

Apologies:

Fires

- OSDF records? - Derek will do some cleanup of the GRACC prod records, continue next week
 - Enabled rollups, reducing the lifetime of OSDF data from TB's to ~3 GB of data.

Announcements

- William on Freshdesk triage next week, Fabio is up the following week
- Ashton is out next week Wed-Fri

Nebraska (Derek, John, Huijun, Ashton)

- **GRACC**
 - Monitoring of the new message bus
 - ~~GRACC Grafana Image renderer to be updated~~
 - Deleting records older than 2 years
- **OASIS**
 - Slowness of underlying Nebraska ceph storage due to migration to different datacenter. We adjusted the rebuild speeds to be less aggressive and are monitoring it.
- **OSG Status Page** (<https://status.osg-htc.org/>)
 -
- **Check_MK - Alerting**
 - Look into upgrading check_mk
 - Tied in with UNL monitoring, will need to coordinate.
 - Remove from check_mk,
 - systemd: cvmfs-data-update@ligo.osgstorage.org.service
 - systemd: cvmfs-data-update@stash.osgstorage.org.service
 -
- **CVMFS origin**
 - Removing Ligo repos
- **XRootD monitoring collector**

-
- **CVMFS Singularity sync**
-
- **Provisioner updates**
 - Ashton put in PR for flux for NRP namespaces (dev, prod)

Madison (Brian Lin, Jeff Peterson, Jason)

- Stay away from `3x-replica-block`, use `3x-replica-hdd-raddus` instead. The former does not get much love and has bad performance + reliability
- Raddus raw NVMe utilization is in better shape by taking manual steps to get the FS to release used blocks (<https://opensciencegrid.atlassian.net/browse/INF-3269>)
- Harbor outage issues: Tempest hosts have been hooked up to a slower switch. UW Madison Harbor instance was in really bad shape so we've been trying to get it up and running

Chicago/Collaborations (Pascal)

- UConn CEs and communication with Richard on their downtime.
 - Checking on the GlueX central manager migration to tiger (from UConn). Successful so far, but waiting for GlueX & Jlab to scale up to production. Note that the GlueX frontend has been on Tiger all along. Need the mapfile issuer to change from Gremlin. Uconn to the new one and distributed to all CEs (HostedCE and non-HostedCE) that allow the GlueXVO. Also note, we have not validated with Jlab flocking to the OSPool yet for that AP.
 - The migration also needs to address the connection to non-JLab APs (so that the new CM can see job pressure there).
 - As a recap - discussed with Jeff D on Thur the addition of new frontends to the tempest factory. Some points were clarified:
 - The UCSD-CMS frontend is only for users with access to UCSD-T2 CMS and the Users Analysis Facility (basically Frank's group, with some users from UCSB)
 - The science is CMS, and pilots go to CMS-T2 institutions and some opportunistic
 - Not connected to the CMS global pool and not available across the CMS community
 - Other options for a factory: CERN? Too far and an overkill. Fermilab? The factory there only submits to the T1.
 - Is it a burden? No, The ucscd-cms frontend was added on June24th but we have seen no significant uptick in the IOPS on tempest there.
- Link:
- https://grafana.tempest.uchicago.edu/d-gtf0Bzik/ceph-pools?orgId=1&refresh=1m&from=now-30d&to=now&viewPanel=4&var-DS_PROMETHEUS=Prometheus&var-interval=30s&var-pool=rbd-3x

Reminder that a redline at 80% capacity stands at 27k. Currently at 18k for Tempest.

We do see an appreciable increase after June 24th in CPU load and memory as expected:

CPU:

<https://grafana.tempest.uchicago.edu/d/85a562078cdf77779eaa1add43ccec1e/kubernetes-compute-resources-namespace-pods?orgId=1&refresh=1d&var-datasource=default&var-cluster=&var-namespace=gfactory-prod&from=now-30d&to=now&viewPanel=5>

Memory:

<https://grafana.tempest.uchicago.edu/d/85a562078cdf77779eaa1add43ccec1e/kubernetes-compute-resources-namespace-pods?orgId=1&refresh=1d&var-datasource=default&var-cluster=&var-namespace=gfactory-prod&from=now-30d&to=now&viewPanel=7>

We should probably check on limits for these two

Michigan (Shawn)

- Waiting to renew the data to transport to Opensearch and FNAL tape

UCSD (Jeff, Fabio)

OSDF:

-

Hosted CEs

- Ashton is working on the Hosted CE Operator
- Migrating glidein shape configs out of factory and into Hosted CEs, completed all we can (94 out of 100); 3 broken, 3 whole node (waiting for GWMS code updates)

GWMS Factory

- Waiting on frontends to start upgrading to 3.11.2
- Planing to test 3.11.2 rc on OSG ITB Factory next week

AOB