Implementing Ceph at GMU's College of Engineering and Computing

Dan O'Brien

Linux, Database & Cloud Administrator

Ceph Days Raleigh

March 25, 2026

George Mason University

- Located in Fairfax, VA
- 10 Colleges and Schools
- 40,378 Students
- 1,529 Instructional Faculty
- 258 Research Faculty
- 1,238 Adjunct Faculty

George Mason University

- Ilia Malinin
- “Quad God”
- Olympic gold medalist

College of Engineering and Computing

- Largest academic unit at George Mason
- 10,624 Students
- 451 Faculty
- $83.7M Research Expenditures (FY 2025)

Central vs. Distributed I.T.

Central I.T.

- **Network services**
  LAN, WiFi, VoIP, Firewall, VPN
- **ID & Authentication**
  AD, LDAP, Kerberos, SSO
- **Data center operations**
- **Overall policy framework**

Office of Research Computing

- **HPC Cluster**
- **OpenStack**
- **Petabyte-scale storage**
  (Another storage vendor)

Distributed I.T.

- **Manage our own infrastructure**
- **Academic software and systems**
- **Computer labs**
- **Desktop support for faculty and staff**
- **Some research support**
- **Facilitates requests to central I.T.**

CEC Infrastructure

HARDWARE	Dell Servers RHEL/Rocky/Ubuntu
NETWORK	10 Gb Ethernet
STORAGE	NFS - Gluster HDFS - Ceph

VIRTUAL SERVERS	OpenStack oVirt
SYSTEMS	Linux Hadoop Oracle MySQL GitLab License Servers

Data Storage Challenges

- Data is not in one place
- Lack consistent redundancy and high availability
- Architecture doesn’t scale well
- Need to support EOL OSes until replaced

- *Ad hoc* management of data lifecycle
- Data lifecycle complicated by adjunct faculty and students leaving and returning

Ceph to the Rescue!

Current Ceph Configuration

Containers with Cephadm

3 storage nodes

2 service nodes

25 OSDs

128 GiB raw storage

HDDs with SSD WAL

10 Gb Ethernet

Storage nodes have 2 x bond

Isolated VLANs for back-end and iSCSI

How CEC is Using Ceph

OpenStack

- First use case for Ceph
- October 2021 - single-node Ceph “cluster”
- Natively supports Ceph as a storage back-end
- Multiple instances using same infrastructure

CEC Storage and Applications

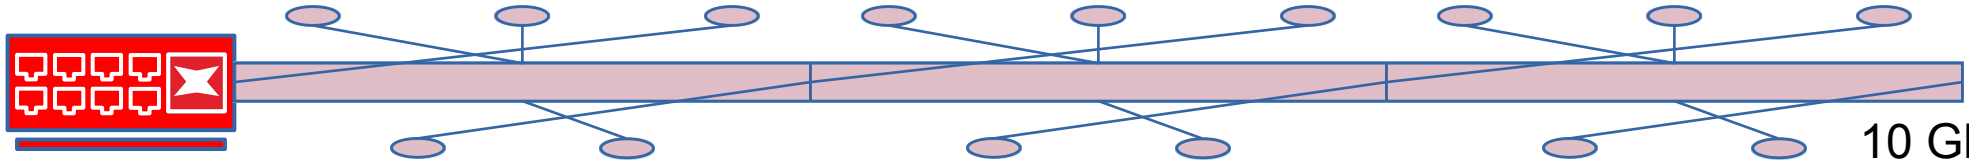

10 Gb Ethernet

Ceph

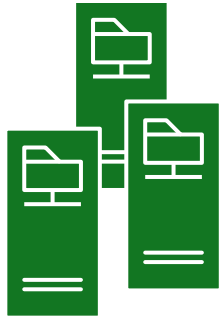

NFS Servers

Gluster

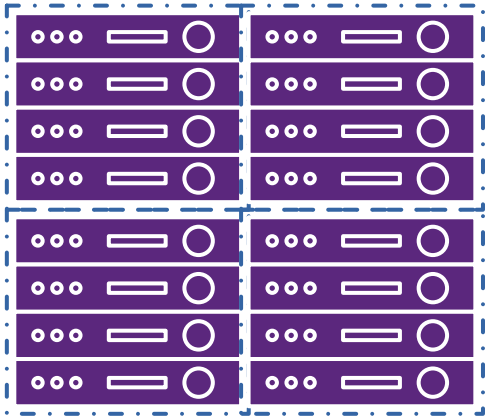

Hadoop

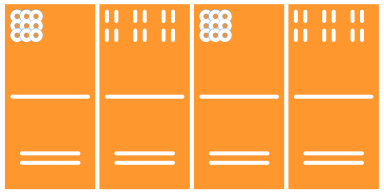

OpenStack

Zeus

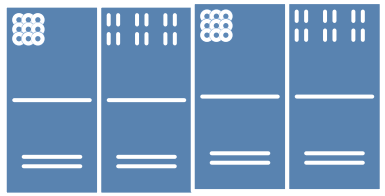

oVirt

How CEC is using Ceph

CephFS for home directories

- Biggest obstacle was understanding our utilization
- Created an inventory tool that enabled us to see usage of home directories across servers
  - User ID, Faculty/Staff or Student, number of files, space used, last updated
  - Integrates user status (faculty/student + active/inactive), and registration data
  - Looks for key properties (incorrect permissions, VScode data)

How CEC is using Ceph

```
sqlite> select * from flags;
flag  flag_desc
----  ------------------------------------------
2     Contains Apache Maven content (.m2)
b     .bash_profile owned by root
c     ~/.cursor-server directory present
C     ~/CS directory present
D     Duplicate homedir - also check archives
E     Current employee
g     Group not idestudent or itefacstaff
G     Group not correct for status (S=500, E=400)
h     helios_html permissions wrong or
          not a directory
j     ~/.cache/JetBrains directory present
m     helios_html missing
o     Go package cache present
p     Permissions incorrect
```

```
flag  flag_desc
----  ------------------------------------------
q     max_files or max_bytes quota is not set
Q     max_files and max_bytes quotas are not set
r     Registered for classes other semester
R     Registered for classes this semester
S     Current student
v     .vscode-server directory present
X     Not a current student or employee
z     Owned by root
```

How CEC is using Ceph

Table: homes

Filter in any column

	hostname	path	netid	uid	gid	perms	modified	num_files	size_kb	last_used	size_mb	status
	ceph	Filter	Filter	Filter	Fil...	Filter	Filter	Filter	Filter	Filter	Filter	Filter
1	ceph	a/v/avarelao	avarelao	672248	500	2700	2024-08-20 16:46	3	2	2024-08-20 16:46	0.002	empty
2	ceph	a/v/avelasq7	avelasq7	672598	500	2700	2026-01-05 09:18	3	2	2026-01-05 09:18	0.002	live
3	ceph	a/v/averman	averman	642799	500	2700	2023-05-24 09:16	4	2	2023-05-24 09:16	0.002	live
4	ceph	a/v/avayalil	avayalil	699095	500	2700	2025-08-18 09:18	3	2	2025-08-18 09:18	0.002	live
5	ceph	a/v/avillal8	avillal8	656134	500	2700	2025-01-10 09:18	3	2	2025-01-10 09:18	0.002	live
6	ceph	a/v/avalenc9	avalenc9	643645	500	2700	2025-05-26 09:21	3	2	2025-05-26 09:21	0.002	live
7	ceph	a/v/avalerio	avalerio	658831	500	2700	2025-08-18 09:18	3	2	2025-08-18 09:18	0.002	live
8	ceph	a/v/avurity	avurity	550443	500	2700	2020-09-08 23:33	5	4	2020-11-19 15:28	0.004	live
9	ceph	a/v/avirk7	avirk7	687643	500	2700	2025-05-26 09:21	3	2	2025-05-26 09:21	0.002	live
10	ceph	a/v/avalluru	avalluru	690418	500	2700	2025-10-12 23:18	56	262	2025-10-28 18:27	0.256	live
11	ceph	a/v/avangeld	avangeld	588309	500	2700	2025-05-06 18:30	8660	68113	2025-05-06 18:30	66.517	live
12	ceph	a/v/avaidya2	avaidya2	663984	500	2700	2024-12-11 18:34	10	61	2024-12-11 18:34	0.06	live
13	ceph	a/v/avasqu36	avasqu36	665349	500	2700	2024-05-21 09:16	4	2	2024-05-21 09:16	0.002	live
14	ceph	a/v/avu24	avu24	686405	500	2700	2025-05-26 09:21	3	2	2025-05-26 09:21	0.002	live
15	ceph	a/v/avasqu38	avasqu38	670161	500	2700	2024-08-20 16:46	3	2	2024-08-20 16:46	0.002	live

How CEC is using Ceph

```
INSERT INTO to_be_moved
SELECT h.netid, h.hostname, h.last_used, h.flags, h.num_files, h.size_mb,
    ( SELECT GROUP_CONCAT(REPLACE(r.course, " ", ""), ":")
        FROM roster r
        WHERE r.netid=h.netid
        GROUP BY r.netid )
    AS courses
FROM homes h
WHERE
    EXISTS ( SELECT course
                FROM roster r
                WHERE r.netid=h.netid
                  AND r.course IN ('CS 112', 'CS 211', 'CS 262', 'CS 367', 'CS 450',
                      'CS 471', 'CS 475', 'CS 531', 'CS 550', 'CS 571', 'CS 657')
            )
    AND status='live'
    AND hostname<>'ceph'
    AND moved IS NULL
ORDER BY hostname, size_mb DESC, netid;
```

How CEC is using Ceph

CephFS for home directories

- A lot of experimentation and planning
- Purchase new hardware and drives
- SSDs for WAL+Db (Thanks, Clyso!)
- Began by creating all new home directories on Ceph (August 2024)
- Change in usage patterns necessitated accelerating process (May 2025)
- Performance problems due to Visual Studio

How CEC is using Ceph

Zeus Cluster Usage Leading Up to Project Due Date

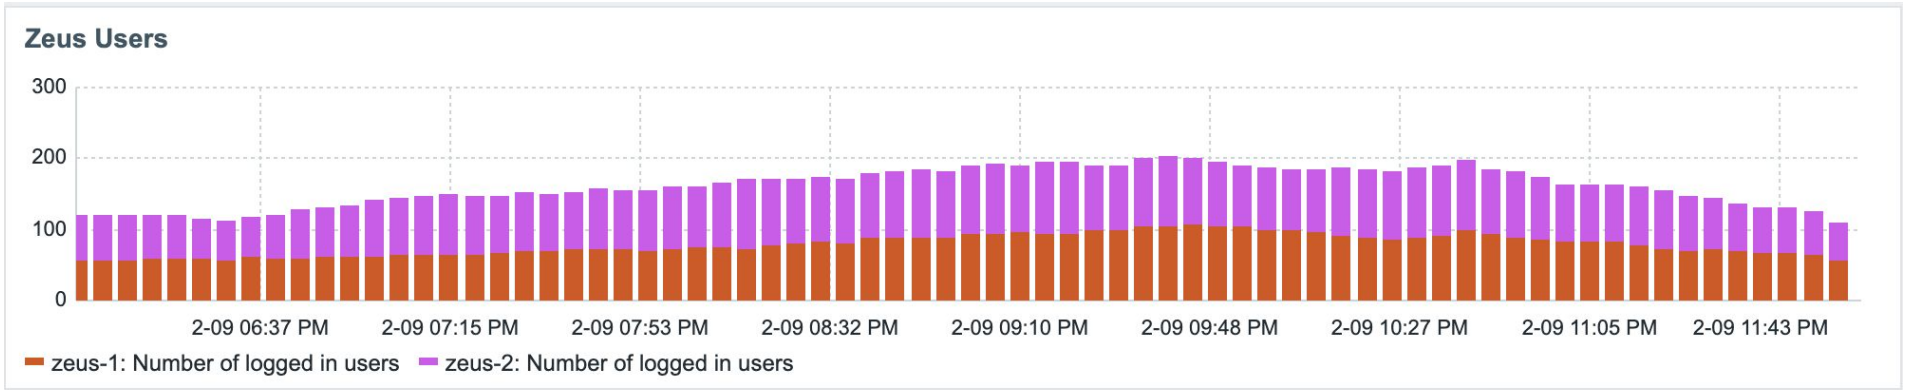

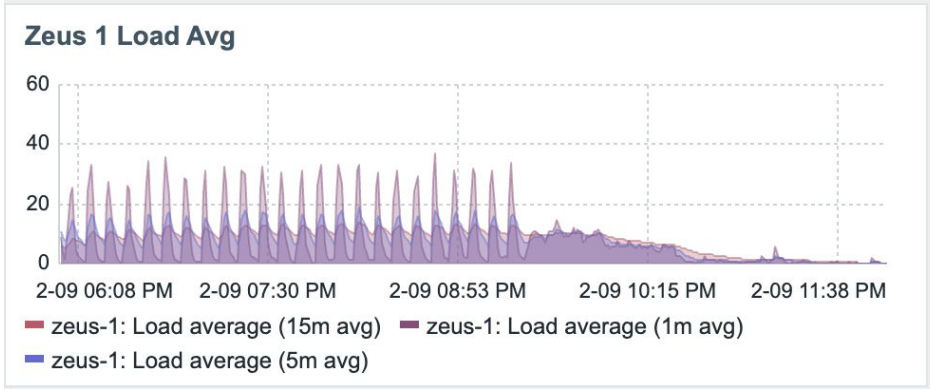

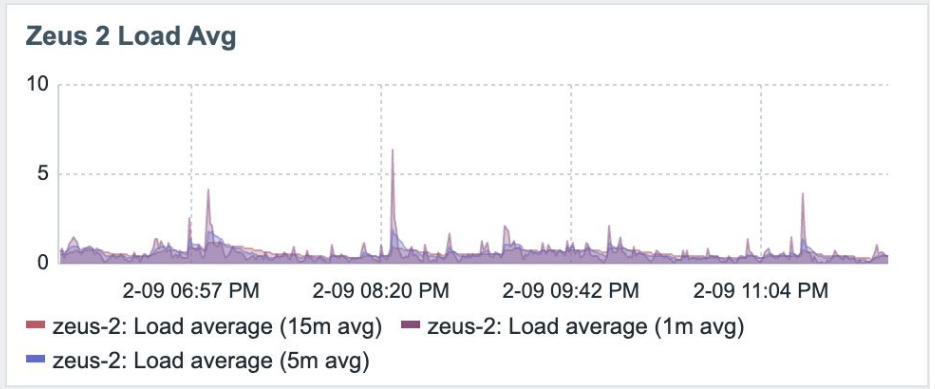

How CEC is using Ceph

CephFS for home directories

- Single CephFS w/ MDS + hot stand-by
- SSDs for WAL+Db after Ceph Day 2024 NYC
- Subvolumes
- Default pool (Replica x 3), plus dedicated pools for other subvols
- Replacing 3 NFS servers
- Current status: 95% complete by number, <50% by data

How CEC is using Ceph

CephFS for home directories

- Script for migration of user data
- Does not migrate low-value data
  - Visual Studio extensions
  - Cache
- Significant space savings due to cleanup and 512b vs 4Kb block size
- Inventory utility allowed for targeted migration of CS students
- No outages (no student extensions) since moving

How CEC is using Ceph

Excluding cache and Visual Studio cruft can reduce data transferred by as much as 90%

```
rsync -av --progress --sparse \
    --exclude .cache/* \
    --exclude .*-server/code-???????????????????????????????????? \
    --exclude .*-server/bin/* \
    --exclude .*-server/cli/servers/* \
    --exclude .*-server/extensions/* \
    --exclude .*-server/data/CachedExtensionVSIXs/* \
    $FROM $DEST/$PREFIX/
```

How CEC is using Ceph

CephFS for home directories

- Quotas
  - Students: 10Gb / 100,000 files, Faculty: 250 GiB / 500,000 files
  - Set during homedir creation or migration
  - Custom scripts to help manage (and request temporary increase)
  - Users see current utilization on logon
  - Inventory process updated to collect data and flag irregularities

How CEC is using Ceph

```
obriend@dans-m4pro-personal ~ % ssh dobrie2@zeus.cec.gmu.edu

Last login: Thu Mar 12 06:37:41 2026 from 10.48.176.125
You have used 3.2Gi (1%) of your 250Gi storage quota
You have used 52365 (10%) of your 500000 files quota

 _____________________________________
/ In charity there is no excess.      \
|                                     |
\ -- Francis Bacon                    /
 -------------------------------------
        \   ^__^
         \  (oo)\_______
            (__)\       )\/\
                ||----w |
                ||     ||
[dobrie2@zeus-2 ~]$
```

How CEC is using Ceph

```
… snip …

FSTYPE=`stat -f -c %T $HOMEDIR`
if [[ $FSTYPE == 'ceph' ]]; then

  CEPH_QUOTA_BYTES=`getfattr -n ceph.quota.max_bytes --only-values $HOMEDIR 2>/dev/null`
  if [[ $CEPH_QUOTA_BYTES == '' ]]; then
    CEPH_QUOTA_BYTES=0
  fi
  if [ $CEPH_QUOTA_BYTES -gt 0 ]; then
    CEPH_USED_BYTES=`getfattr -n ceph.dir.rbytes --only-values $HOMEDIR 2>/dev/null`
    PCT_STORAGE=$(( $CEPH_USED_BYTES * 100 / $CEPH_QUOTA_BYTES ))
    stoplight $PCT_STORAGE
    echo -n "You have used `echo $CEPH_USED_BYTES | /usr/bin/numfmt --to iec-i` ($PCT_STORAGE%) "
    echo -n "of your `echo $CEPH_QUOTA_BYTES | /usr/bin/numfmt --to iec-i` storage quota"
    echo $COLOR_NORMAL
  else
    echo "You do not have a storage quota"
  fi

… snip …

fi
```

How CEC is using Ceph

NFS Gateway

- Primarily used for compatibility for EOL operating systems
- Also for instances where systems don’t have direct access to storage network
- HAProxy+Keepalived is currently broken (Reef)

How CEC is using Ceph

Consolidation of web content

- Having CephFS made this an easy step
- Subvolume + dedicated pool allows segregation for improved security
- Decoupling content storage and delivery
  - Better visibility into what we're hosting
  - Consolidated backup
  - Updates and capacity expansion become easier

How CEC is using Ceph

oVirt Upgrade

- OpenStack available, but new oVirt cluster provided a simpler migration path - quicker retirement of EOL hosts in old cluster
- Ceph iSCSI gateway (oVirt self-hosted engine only supports iSCSI), but exploring “Managed Block Storage” (Ceph)
- Better support for current network configuration
- Trunk ports w/ VLAN tagging
- Direct access to the storage network

How CEC is using Ceph

Object Storage

- March 2024
- When used with OpenStack it *replaces* Swift
- Can also be used directly with S3-compatible API
- Would like to use this for archival storage of home directories, or as an alternative to large static data in homedir

What We've Achieved So Far

BEFORE

- Data is not in one place
- Lack consistent redundancy and high availability
- Architecture doesn't scale well

AFTER

- Single storage platform for multiple applications
- Highly available and resiliant
- Easily scales as storage needs evolve

The Future

- **Complete migration of NFS data**
  - NFS will continue to be used for backups and “cold” storage
- **Build out Ceph cluster**
  - Reef → Squid → Tentacle
  - Add dedicated SSD storage tier
  - Full(er) separation of storage and service nodes
  - Erasure coding

The Future

- **Active-Active MDS w/ directory pinning**
- **Scheduled snapshots for CephFS**
  - Increased granularity for file recovery (currently only daily)
  - Better recovery for lost student work and academic integrity
- **Pool replication for business continuity**

The Future

- **Storage-as-a-Service for faculty**
  - Eliminate *ad hoc* storage (NAS devices)
  - Compliance with university policies and DUAs
  - Object storage as an option
  - Allow for usage of internal resources instead of public cloud
- **Maybe better Windows/AD integration with Samba?**

GEORGE MASON UNIVERSITY

Questions?

Implementing Ceph at GMU's College of Engineering and Computing

Dan O'Brien

Linux, Database & Cloud Administrator

Ceph Days Raleigh

March 25, 2026