

OSiRIS: A Distributed Storage and Networking Project Update



Open Storage Research Infrastructure

Shawn McKee for the OSiRIS Collaboration

University of Michigan, Indiana University, Michigan State University,
Wayne State University

UM Storage Community of Practice (CoP)

April 29, 2020

Introduction



Today I want to provide an update on the **OSiRIS** project, a 5-year, \$5M storage infrastructure project, lead by U-M.

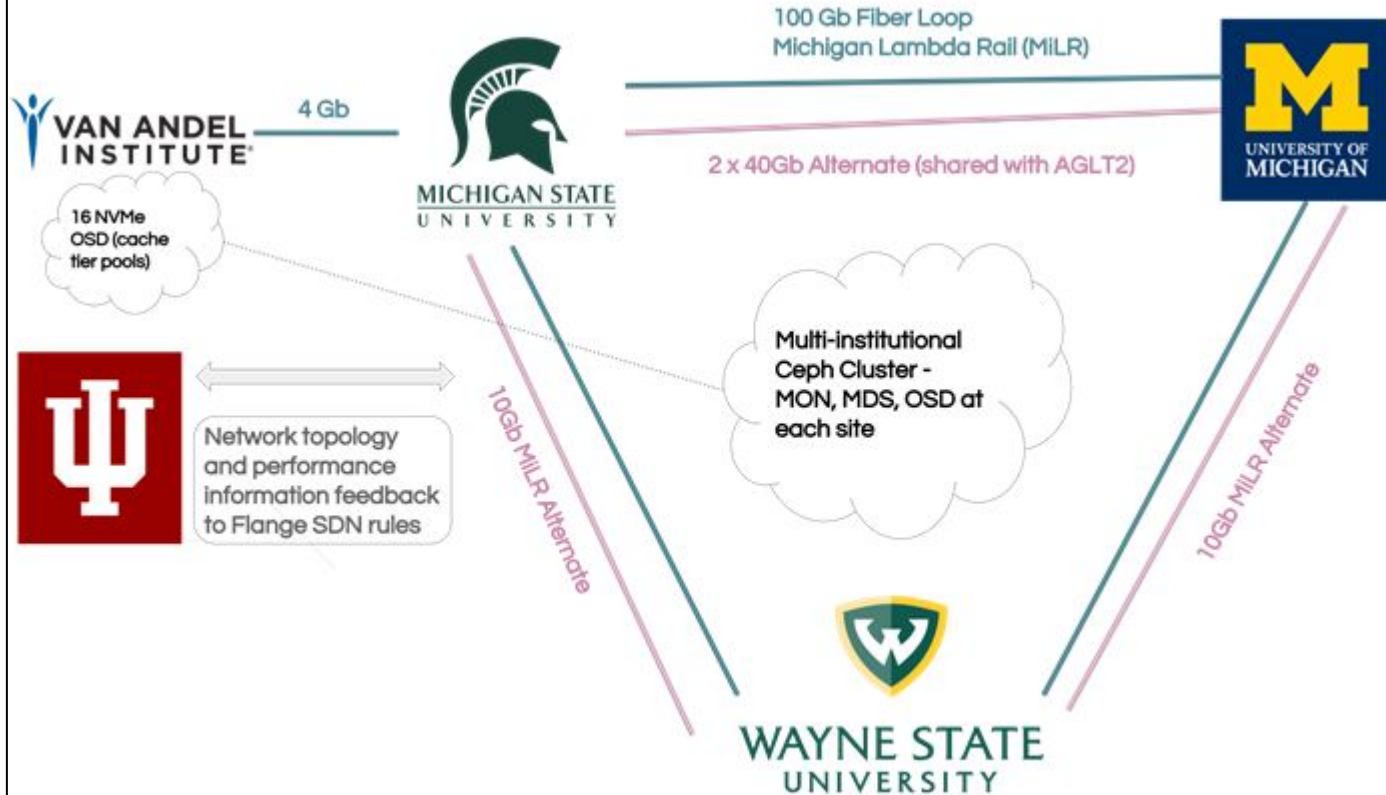
I wasn't clear how many of you would already have seen presentations about the project, so I will present a mix of overview and update.

Please feel free to ask questions or inject comments as we go.

OSiRIS Overview (Review)

The OSiRIS proposal targeted the creation of *a distributed storage infrastructure, built with inexpensive commercial off-the-shelf (COTS) hardware, combining the Ceph storage system with software defined networking to deliver a scalable infrastructure to support multi-institutional science.*

Current: Single Ceph cluster (**Nautilus 14.2.4**) spanning **U-M, WSU, MSU** - **1368 OSD / 13.7 PiB**



OSiRIS Storage Summary

We have **deployed 13.7 pebibytes (PiB) of raw Ceph storage** across our **three research institutions** in the state of Michigan.

- Older storage node is a 2U headnode and SAS attached 60 disk 5U shelf with either 8 TB or 10 TB disks, 4x25G network links (two dual 25G cards)
- New **year-4** hardware installed
 - Dell R7425 (Dual AMD 7301) 2U, 16x12TB disks, 128G RAM, 2x25G NIC, 2x10G, 1x1G, 4 x Samsung 970 Pro 512G NVMe, BOSS card
 - Added **6.3 PiB** to OSiRIS by January 2020 (now 1368 total disks)
- Ceph components and services are virtualized

The **OSiRIS hardware is monitored by Prometheus** and configuration control is provided by Puppet

Institutional identities are used to authenticate users and authorize their access via CoManage and Grouper

Augmented **perfSONAR** is used to **monitor and discover the networks** interconnecting our main science users.

The primary driver for OSiRIS was a set of science domains with either **big data** or **multi-institutional** challenges.

OSiRIS is supporting the following science domains:

- **ATLAS** (high-energy physics), **Bioinformatics**, **Jetscape** (nuclear physics), **Physical Ocean Modeling**, **Social Science** (via the Institute for Social Research), **Molecular Biology**, **Microscopy**, **Imaging & Cytometry Resources**, **Global Night-time Imaging**
- We are currently “on-boarding” new groups in **Genomics and Evolution** and **Neural Imaging** (next slide)
- **Primary use-case is sharing working-access to data**

Recent Science Domains

Brainlife.io (Neuroimaging) - Brainlife organizes neuroimaging data and data derivatives using their registered data types. No single computing resources has enough storage capacity to store all datasets, nor reliable enough so that user can access the data when they need them. They will depend on OSiRIS to store datasets and transfer data between computing resources.



Oakland University - Already a user of MSU iCER compute resources, OU will leverage OSiRIS to bring their data closer for analysis and for collaboration with other institutions.

Evolution - Large-scale evolutionary analyses, primarily phylogenetic trees, molecular clocks, and pangenome analyses

Genomics - High volume of human, mammal, environmental, and intermediate analysis data

New and Ongoing Collaborations

Open Storage Network - We will be providing ~1 PB to be included in the Open Storage Network (<https://www.openstoragenetwork.org>)

- Timeline depends on OSN readiness to engage, some discussions at OSN group meeting at TACC in Fall 2019

FABRIC - This is a newly funded NSF project to create a network testbed at-scale (1.2 Tbps across the US). OSiRIS will be an early adopter/collaborator, providing ~1 PB to support science use-cases

Library Sciences - OSiRIS roadmap plans for data lifecycle mgmt

- Following detailed analysis of two specific datasets, library scientists at UM are working on automated metadata capture and indexing
- Integration with U-M 'Deep Blue Data' archival system also planned

OSiRIS News for 2020

The start of 2020 saw some **major changes** for the OSiRIS project.

Ben Meekhof (who many of you know) has been the primary OSiRIS engineer since the project was started. Ben took a great opportunity to join Eric Boyd's networking team at the beginning of January and we have reorganized to try to fill in Ben's role in the project

- The lead engineers at MSU (Andy Keen), WSU (Michael Thompson) and a new OSiRIS hire from Fall 2019 (Soundar Rajendran) jointly cover Ben's previous role.

In February 2020, **Ezra Kissel / Indiana**, who has led the NMAL work for OSiRIS, took a job with ESnet (Energy Sciences network).

- **Jeremy Musser / Indiana**, a long-time graduate student on the project has been filling in for Ezra and we are working on engaging additional effort from IU

OSiRIS is in its 5th year of a 5 year grant but in **March 2020** we have successfully requested and received a **no-cost extension till the end of August 2021**.

MSU has found a suitable hire who will join OSiRIS at 50% time starting in May.

Network Upgrades - 100Gb MiLR



MiLR is a **high-speed, special purpose, data network** built jointly by Michigan State University, the University of Michigan, and Wayne State University, and operated by the Merit Network.

Thanks to combined effort from campus network teams and Merit we were able to deploy direct 100Gb links via MiLR fiber landing directly on our OSiRIS rack switches

- Now we have more options for network management without campus network disruptions and this provides options for experimenting with SDN via NMAL

In our current phase of implementation, they carry only the Ceph 'cluster network' used for OSD replication data (Ceph self-healing)

Normal ceph recovery/backfill operations could easily overwhelm smaller links with this traffic, so removing it was a **huge difference** that let us completely remove throttles on Ceph recovery (see next slide)

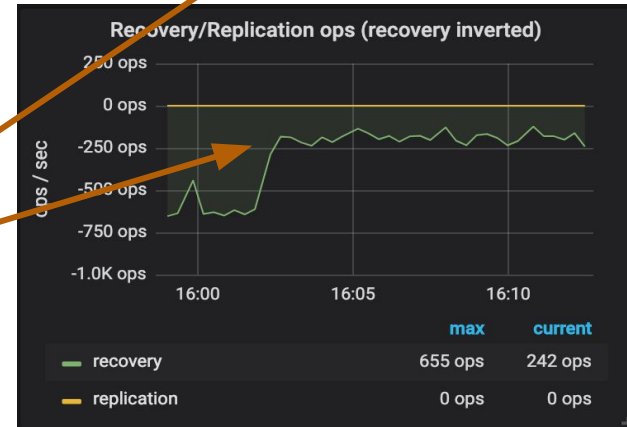
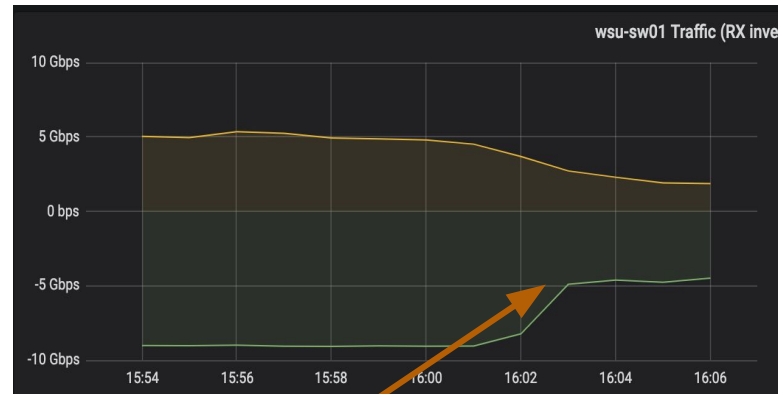
Unbalanced Networks and Ceph

Prior to our installation of 100G links for Ceph cluster backend we had issues with network bandwidth inequality: U-M and MSU sites had 80G link to each other but 10G to WSU datacenter

- Adding a new node, or losing enough disks, would completely swamp the 10G link and cause OSD flapping, mon/mds problems, service disruptions

Lowering recovery tunings fixed the issue, at the expense of under-utilizing our faster links. Recovery sleep had the most effect, the others not as clear

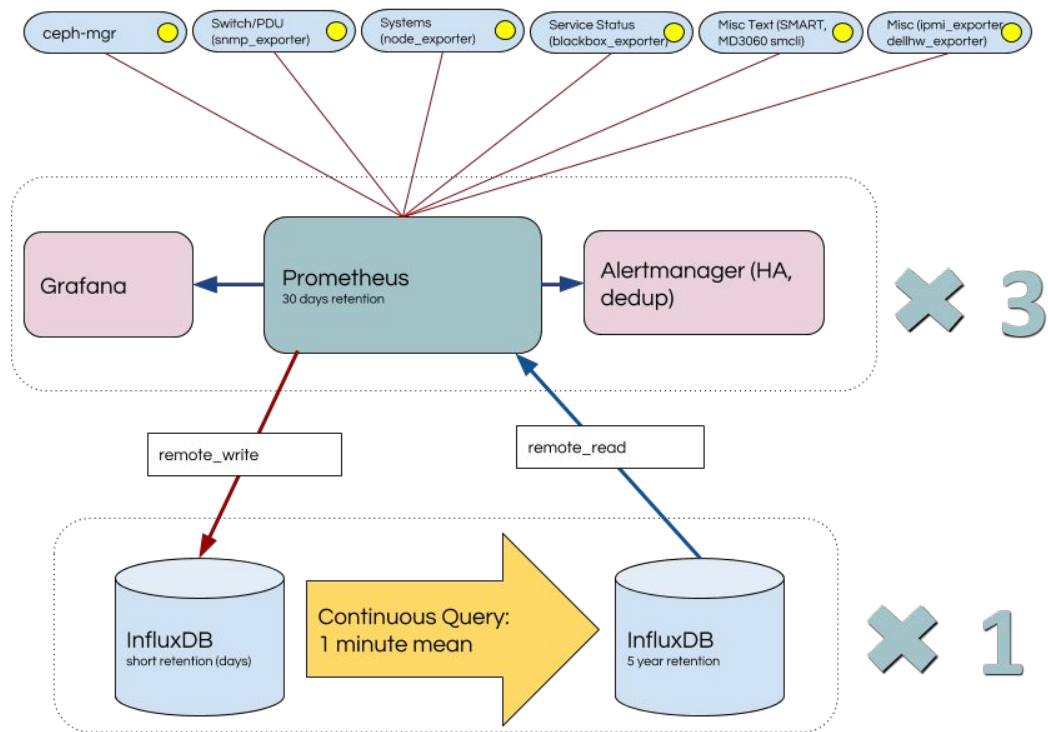
```
osd_recovery_max_active: 1      # (default 3)
osd_backfill_scan_min: 8        #(def 64)
osd_backfill_scan_max: 64      #(def 512)
osd_recovery_sleep_hybrid: 0.1  #(def .025)
```



Monitoring and Metrics with Prometheus

Recently we consolidated all of our metrics, monitoring, alerting to Prometheus

- Migrated from a combination of Influxdb and Collectd
- Continue to use Grafana to visualize, Influxdb for long-term retention
- Consideration was given to standing up more of the influx (TICK) stack, pros and cons each way
- Text collector scripts and alert rules in our git repo (grafana dashboards soon)



<https://github.com/MI-OSiRIS/osiris-monitoring>

CheckMK Service Monitoring (MSU, U-M, WSU)

msu-omd										msu-gw01										msu-idap01									
state	Host	Icons	OK	Wa	Un	Cr	Pd	state	Host	Icons	OK	Wa	Un	Cr	Pd	state	Host	Icons	OK	Wa	Un	Cr	Pd						
UP	lu-omd		21	1	0	0	0	UP	msu-gw01		1	0	0	0	0	UP	msu-idap01		32	1	0	0	0						
UP	msu-mono01		28	1	0	0	0	UP	msu-prov		29	2	0	0	0	UP	msu-prov		25	2	0	0	0						
UP	msu-ps01		30	3	2	0	0	UP	msu-sdn01		28	1	2	0	0	UP	msu-stor01		100	2	0	0	0						
UP	msu-stor02		99	2	1	0	0	UP	msu-stor03		99	2	1	0	0	UP	msu-stor04		99	2	1	0	0						
UP	msu-stor05		100	2	0	0	0	UP	msu-sw01		26	0	0	0	0	UP	msu-sw02		37	0	0	0	0						
UP	msu-virt01		47	1	0	0	0	UP	rac-msu-ps01		42	1	0	0	0	UP	rac-msu-stor01		63	1	2	0	0						
UP	rac-msu-stor02		65	1	4	0	0	UP	rac-msu-stor03		69	1	0	0	0	UP	rac-msu-stor04		56	8	4	0	0						
UP	rac-msu-stor05		50	1	0	0	0	UP	rac-msu-virt01		51	0	0	0	0	UP	um-omd		31	2	0	0	0						
UP	wsu-omd		29	2	0	0	0																						

um-omd										msu-omd										msu-omd									
state	Host	Icons	OK	Wa	Un	Cr	Pd	state	Host	Icons	OK	Wa	Un	Cr	Pd	state	Host	Icons	OK	Wa	Un	Cr	Pd						
UP	comanager		27	1	0	0	0	UP	lu-omd		24	0	0	0	0	UP	msu-omd		31	1	0	0	0						
UP	oproc-agt2-orig		34	1	0	0	0	UP	rac-um-grafana		6	0	0	0	0	UP	rac-um-ps01		64	0	0	0	0						
UP	rac-um-stor01		81	0	0	0	0	UP	rac-um-stor02		84	0	0	0	0	UP	rac-um-stor03		87	0	1	0	0						
UP	rac-um-stor04		1	1	0	0	0	UP	rac-um-stor05		73	0	0	0	0	UP	rac-um-stor06		66	0	0	0	0						
UP	rac-um-virt01		76	0	0	0	0	UP	rac-um-virt01		64	0	0	1	0	UP	um-admin01		28	1	0	0	0						
UP	um-ek		28	2	0	0	0	UP	um-grafana		60	1	0	0	0	UP	um-idap01		31	1	0	0	0						
UP	um-mono01		27	1	0	0	0	UP	um-omd		33	1	0	0	0	UP	um-ps01		50	2	0	0	0						
UP	um-puppet		29	1	0	0	0	UP	um-puppetlb		30	1	0	0	0	UP	um-repo		28	1	0	0	0						
UP	um-sdn01		30	0	0	0	0	UP	um-stor01		141	1	0	0	0	UP	um-stor02		139	1	0	0	0						
UP	um-stor03		140	1	0	0	0	UP	um-stor04		139	1	0	0	0	UP	um-stor05-m		135	1	0	0	0						
UP	um-stor06-m		133	1	0	0	0	UP	um-stor-ds01		83	1	0	0	0	UP	um-stor-ds02		83	1	0	0	0						
UP	um-stor-ds03		83	1	0	0	0	UP	um-stor-ds04		83	1	0	0	0	UP	um-stor-ds05		83	1	0	0	0						
UP	um-stor-ds06		83	1	0	0	0	UP	um-stor-ds07		83	1	0	0	0	UP	um-stor-ds08		83	1	0	0	0						
UP	um-stor-ds09		83	1	0	0	0	UP	um-stor-ds10		83	1	0	0	0	UP	um-stor-ds11		83	1	0	0	0						
UP	um-stor-hvmd01		38	0	0	0	0	UP	um-sv01		43	0	0	0	0	UP	um-sv02		36	0	0	0	0						
UP	um-sv03		25	0	0	0	0	UP	um-virt01-m		68	0	0	0	0	UP	um-virt02		47	2	0	0	0						
UP	um-vler01		52	1	0	0	0	UP	umag12-ps01		31	0	0	0	0	UP	wiki		32	0	0	0	0						
UP	wsu-omd		31	1	0	0	0																						

wsu-omd										msu-omd										msu-omd									
state	Host	Icons	OK	Wa	Un	Cr	Pd	state	Host	Icons	OK	Wa	Un	Cr	Pd	state	Host	Icons	OK	Wa	Un	Cr	Pd						
UP	lu-omd		24	0	0	0	0	UP	msu-omd		31	1	0	0	0	UP	rac-wsu-ps01		1	0	0	0	0						
UP	rac-wsu-stor01		79	0	0	0	0	UP	rac-wsu-stor02		69	0	0	0	0	UP	rac-wsu-stor03		85	0	0	0	0						
UP	rac-wsu-stor04		64	0	0	0	0	UP	rac-wsu-stor05		64	0	0	0	0	UP	rac-wsu-virt01		47	0	0	0	0						
UP	rac-wsu-virt02		57	0	0	2	0	UP	um-omd		33	1	0	0	0	UP	um-ek		27	0	0	0	0						
UP	wsu-grafana		29	0	0	0	0	UP	wsu-idap01		32	0	0	0	1	UP	um-mono01		29	0	0	0	1						
UP	wsu-omd		29	1	0	0	0	UP	wsu-prov		0	0	0	0	0	UP	wsu-ps01		0	0	0	0	0						
UP	wsu-sdn01		30	0	0	0	0	UP	wsu-stor01		143	0	0	0	0	UP	wsu-stor02		136	0	0	0	0						
UP	wsu-stor03		141	0	0	5	0	UP	wsu-stor04		131	0	0	0	0	UP	wsu-stor05		132	0	0	0	0						
UP	wsu-stor-ds01		91	0	0	0	0	UP	wsu-stor-ds02		81	0	0	0	0	UP	wsu-stor-ds03		81	0	0	0	0						
UP	wsu-stor-ds04		80	0	0	0	0	UP	wsu-stor-ds05		79	0	0	0	0	UP	wsu-stor-ds06		81	0	0	0	0						
UP	wsu-stor-ds07		80	0	0	0	0	UP	wsu-stor-ds08		81	0	0	0	0	UP	wsu-stor-ds09		80	0	0	0	0						
UP	wsu-stor-ds10		80	0	0	0	0	UP	wsu-stor-ds11		80	0	0	0	0	UP	wsu-sw01		18	0	0	0	0						
UP	wsu-sw02		54	0	0	0	0	UP	wsu-virt01		103	1	0	0	0	DOWN	wsu-virt02		98	0	0	0	6						
UP	wsu-vler01		91	0	0	1	0																						

OSIRIS - Open Storage Research Infrastructure

refresh: 30 secs

CManage Credential Management

Home > OSiRIS > Ceph Credentials

Ceph Credentials

Description	User ID	Credentials
S3/RGW Access Key	bmeekhof	access_key: 4R7GOZ3XRFS807F3D9WM secret_key: 9sFmGDa63t1ox
Enter New Id:	<input type="text" value="bad%#\$example"/>	<input type="button" value="+"/> Can only contain alphanumeric
S3/RGW Access Key	myservice	access_key: 05RBS3I78HP7U secret_key: 6rKmKoomRzjG
S3/RGW Access Key	myservice	access_key: C136YO80W76T secret_key: 68O33P10Salg2
Ceph Client Key	osiris.bmeekhof	[client.osiris.bmeekhof] key = AQACon1cZSs4FBAA

Add a new userid

Add additional access/
secret pair to this userid

Remove this id and all
keypairs associated

CManage Ceph Provisioner plugin provides user interface to manage S3 credentials and default bucket placement

Work is underway to include a full GUI for managing buckets: Create, rename, download, set ACL from OSiRIS groups or specific user, etc.

Click here to open
selector and choose
new bucket placement

- ✓ OsirisAdmin
- ATLAS
- ARCTS
- osiris.bmeekhof

S3 Fuse Client Bundle

Technically S3 storage makes more sense for most use cases wanting to compute with OSiRIS storage from campus or off-campus locations

- But...not everyone is very familiar with S3
- People often think we are telling them to go use Amazon just by saying S3

We try to make it a little easier by putting together a bundle that automatically FUSE mounts their S3 buckets with s3fs-fuse utility

- Includes setup script, user plugs in credentials
- Auto-detects which OSiRIS S3 endpoint URL is reachable and passes to mount command (our campus cluster users may only be able to reach on-campus endpoint)
- Includes build of s3fs-fuse util made with appimage to be portable to any Linux system.
- <https://github.com/MI-OSiRIS/osiris-bundle>
- <http://www.osris.org/documentation/s3fuse.html>

Globus and Gridmap

We provide Globus access to CephFS and S3 storage

- Ceph connector uses radosgw admin API to lookup user credentials and connect to endpoint URL with them

Credentials: CILogon + globus-gridmap

- CILogon DN in LDAP voPerson CoPersonCertificateDN attribute

We wrote a Gridmap plugin to lookup DN directly from LDAP (student project)

- <http://myumi.ch/lxW1Z>
- <http://myumi.ch/pd0be>

Having the subject DN and lookup entirely in LDAP means it will be easy to add capabilities to COmanage so users can self-manage this information

- Users already self-manage SSH login keys in COmanage (also in LDAP)

Network Management

The OSiRIS [Network Management Abstraction Layer](#) is a key part of the project with several important focuses:

- [Capturing site topology and routing information](#) from multiple sources: SNMP, LLDP, sflow, SDN controllers, and existing topology and looking glass services
- Converge on common [scheduled measurement architecture](#) with existing **perfSONAR** mesh configurations
- [Correlate long-term performance measurements](#) with passive metrics collected via other monitoring infrastructure

Recently wrote new Prometheus exporter to collect perfSonar test results from central ESmond store for alerting and visualization

We demo'ed SDN architecture for traffic routing and traffic shaping / QOS (prioritize client / cluster service traffic over recovery) at **SC19**



OSiRIS Topology Discovery and Monitoring

UNIS-Runtime release integrated into ZOF-based discovery app

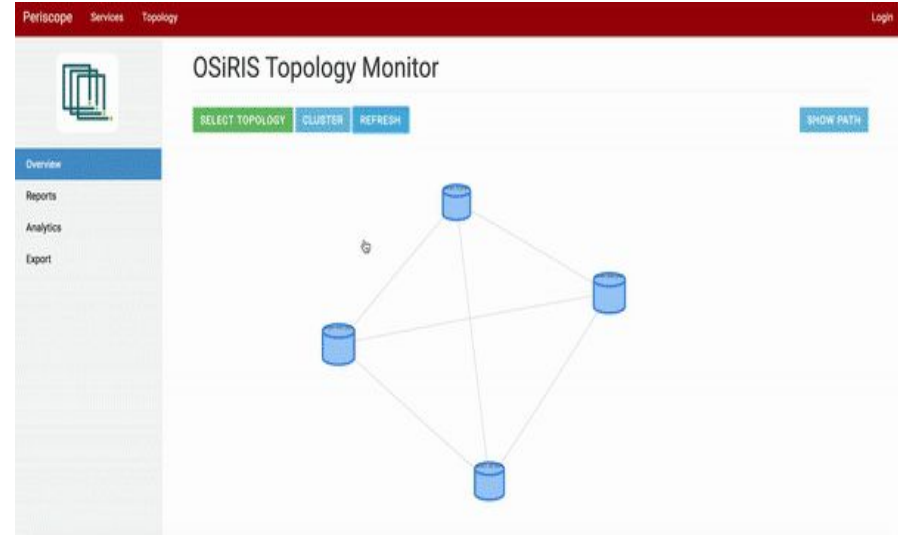
- Increased stability and ease of deployment
- Added extensions for Traceroute and SNMP polling

Web Development has focused on bringing measurements to dashboard

- Link and node highlighting with thresholds determined by link capacities
- Overlay for regular testing results to bring “at-a-glance” diagnostics

Filtering to show layer-2 topology versus layer-3 and virtualized components

- Fault localization, clustering, and zoom are work-in-progress



Summary

OSiRIS continues to improve on our user experience and engage with new collaborators

- ATLAS/High-energy Physics has been a long time user for Event Service data

Our new hardware purchases this year will increase our node count and make EC pools feasible (no clients yet for this)

We look forward to participating in more national scale projects such as the Open Storage Network, FABRIC, Eastern Research Network

On our roadmap this year:

- Upgrade to Octopus (15.2.1)
- Make our S3 services more highly available with LVS failover endpoints on each campus
- Make S3 services more performant by greatly increasing instance count behind the proxy endpoints
- Improve user GUI for managing storage access
- Build more convenient client bundles, modules, etc to make OSiRIS usage as easy as possible
- Adding ATLAS dCache storage to explore using Ceph to manage back-end storage.

Acknowledgements

We would like to thank our **OSiRIS science partners** and our host institutions for their contributions to work described.

In addition we want to explicitly acknowledge the support of the **National Science Foundation** which supported this work via:

- **OSIRIS grant, NSF OAC-1541335**

Questions?

Email us with questions: osiris-help@umich.edu

Website: <http://www.osris.org/>

Further Information

OSiRIS

<http://www.osris.org> project website

Details in various presentations at <http://www.osris.org/publications>

IRIS-HEP

<https://iris-hep.org/> project website

Details in various presentations at <https://iris-hep.org/presentations/bymonth>

DOMA

<https://iris-hep.org/doma.html> sub-project website

DOMA Presentations are available at the above URL

Some Caching studies

https://indico.cern.ch/event/770307/contributions/3301625/attachments/1807559/2952167/Scheduling_with_Virtual_Placement_for_Site_Jamboree.pdf

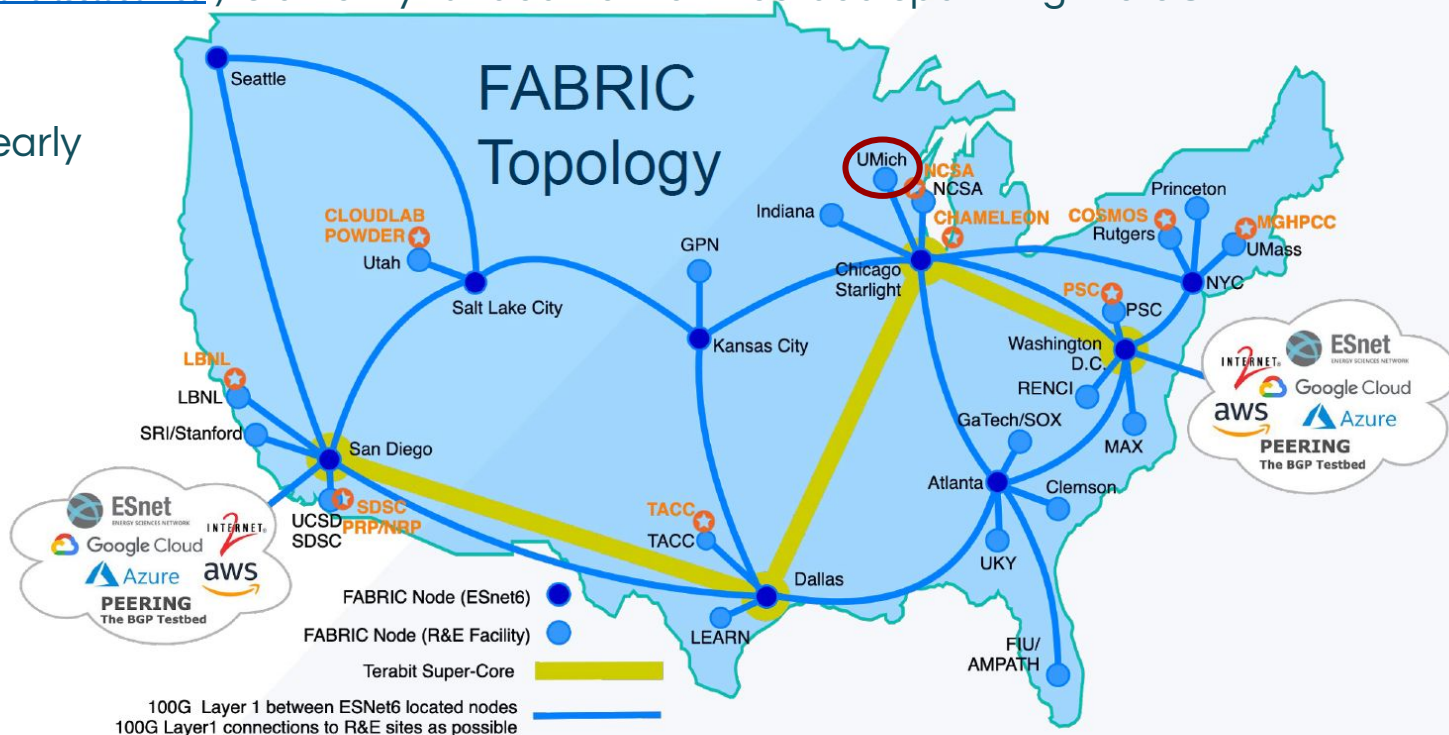
Additional Slides Follow

Backup Slides

FABRIC Topology

FABRIC (<https://fabric-testbed.net/>) is a newly funded network testbed spanning the US

Michigan is an early adopter (2021)



<https://whatisfabric.net/events/fabric-community-workshop-2020>



CManage - Virtual Org Provisioning

When we create CManage COU (virtual org):

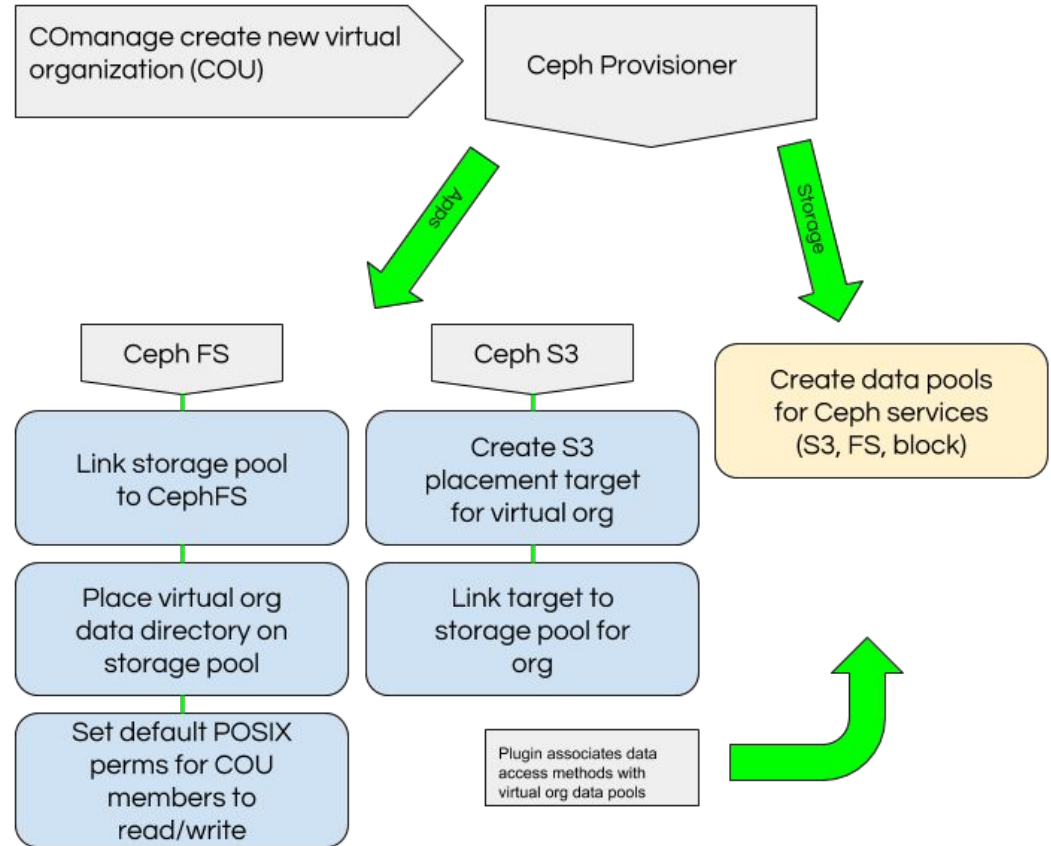
Data pools created

RGW placement target defined to link to pool
cou.Name.rgw

CephFS pool create and added to fs

COU directory created and placed on CephFS pool

Default perms/ownership set to COU all members group, write perms for admins group (as a default, can be modified)



Ideal Facilities

If we could have our way, we would have **ideal facilities**:

- CPUs would always be busy running science workflows
- Any data required would always be immediately available to the CPU when needed
- (Oh, and the facilities would be free and self-maintaining and use negligible power!)

As we all know, it is hard to create efficient infrastructures that manage access to large or distributed data effectively

Approaching “ideal” becomes very expensive (in \$’s and effort)

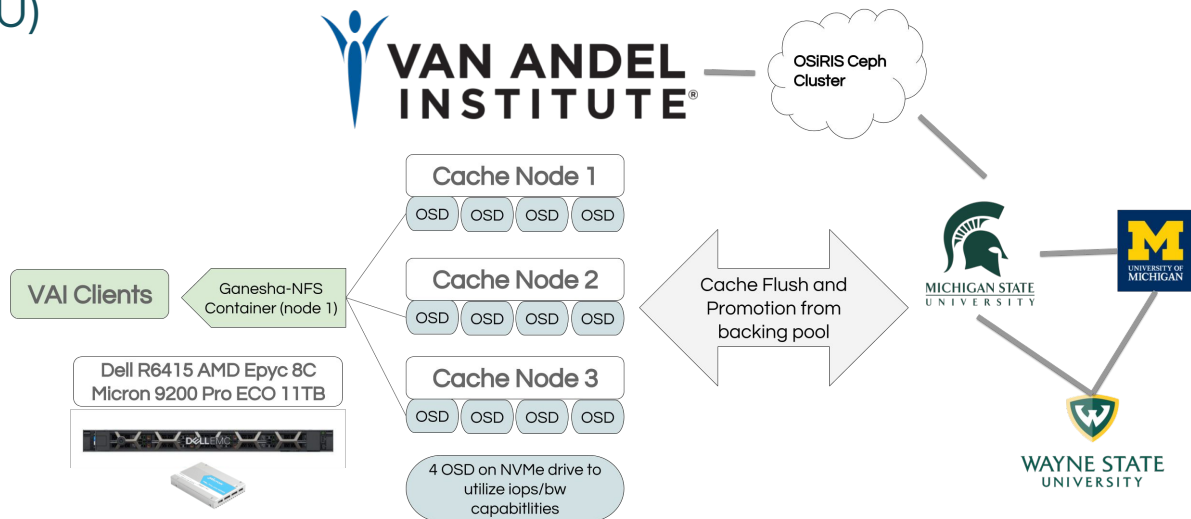
So we need to make progress as best we can.

OSiRIS Ceph Cache Tiers

At Supercomputing conferences (2016/17/18) we've experimented with Ceph cache tiering to work around higher latency to core storage sites

- Deploy smaller edge storage elements which intercept reads/writes and flush or promote from backing storage as needed

Have edge OSiRIS site leveraging this technique at Van Andel Institute (primarily led by MSU)



OSiRIS: Quality of Service for Ceph

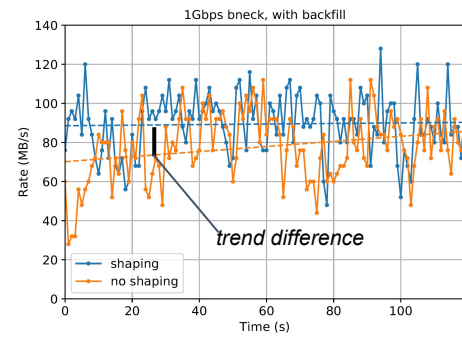
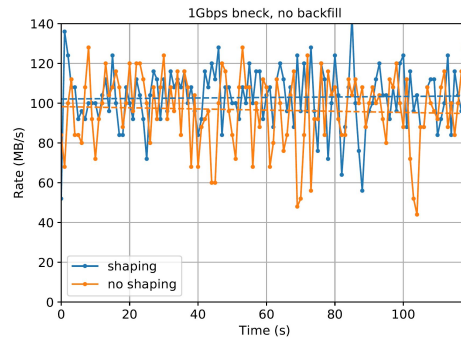
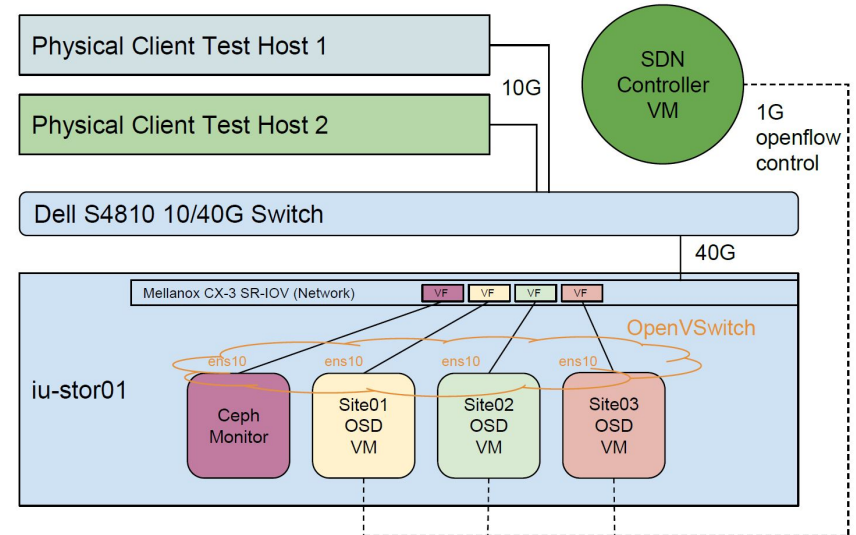
Testbed created to develop QoS functionality

- Explicit control of operations, no noise
- Reduce risk of breaking production

Apply priority queues to ensure that adequate bandwidth exists for Ceph client operations to prevent timeouts and delayed read/write performance

Apply traffic shaping to provide better transport protocol performance between sites with asymmetric link capacities. This is of particular importance when latency between sites is increased

Preliminary results: shaping from sites towards bottleneck can improve client performance, approx 5-10% in early testing.



OSiRIS Lesson's Learned

OSiRIS works very well on a regional scale (networking RTT $\sim < 10$ ms)

We explored scaling for a single Ceph cluster at SC16 where we dynamically added a new site on the exhibition floor 42 ms RTT from the rest of OSiRIS

- The benchmark work-flow data access dropped from **1.2 GB/sec** to **0.45 GB/sec**
- The infrastructure continued to work without problems

Using 'netem' we were able to programmatically add arbitrary delay into the network stack of one of our Ceph servers.

- As we increased the latency we saw the expected impact in throughput
- When we reached **160 ms**, our (untuned) Ceph cluster stopped working
- We needed to decrease the latency back down to **80 ms** to recover

To reach more distributed deployments, OSiRIS would need to start using Ceph Federations (with associated costs) or employ caching to "hide" the latency as much as possible. In DOMA terms, OSiRIS would be appropriate as an element of a data lake.