



AGLT2 Site Report

Shawn McKee/University of Michigan

Bob Ball, Chip Brock, Philippe Laurens, Ben
Meekhof, Ryan Sylvester, Richard Drake

OSG AHM 2016 / Clemson



Outline

-  Site Summary
-  Personnel Changes
-  Moving locations at UM?
-  Site Details

Site Summary

- ❑ The ATLAS Great Lake Tier-2 (AGLT2) is a distributed LHC Tier-2 for ATLAS spanning between UM/Ann Arbor and MSU/East Lansing. Roughly 50% of storage and compute at each site
 - ❑ 7116 single core job slots (Adding 1104 slots (23xR630s) in one month; E5-2680v3, 192 GB ram, 4x500G NL-SAS each)
 - ❑ MCORE slots 550 (dynamic) + 10 (static)
 - ❑ 720 Tier-3 job slots usable by Tier-2
 - ❑ Average 9.65 HS06/slot
 - ❑ 3.7 Petabytes of storage
 - ❑ Total of 68.6 kHS06
 - ❑ Most Tier-2 services virtualized in VMware
- ❑ 2x40 Gb inter-site connectivity, UM has 100G to WAN, MSU has 10G to WAN, lots of 10Gb internal ports and 16 x 40Gb ports
- ❑ High capacity storage systems have 2 x 10Gb bonded links
- ❑ 40Gb link between Tier-2 and Tier-3 physical locations

Personnel Changes

- ❏ After many years with AGLT2 **Ben Meekhof** is moving on. Fortunately for us it is only across campus.
 - ❏ Ben is the new OSiRIS Project Engineer (See talks today/tomorrow).
 - ❏ Many, many thanks to Ben for all his hard work and innovations for AGLT2. He will be missed.
- ❏ We interviewed 9 candidates as a replacement (out of about 18 applications) and selected **Ryan Sylvester** as the best candidate.
- ❏ Ryan will be starting in approximately 1 week and we hope he will be able to attend our next Facilities get together in person.

Possible Relocation at UM

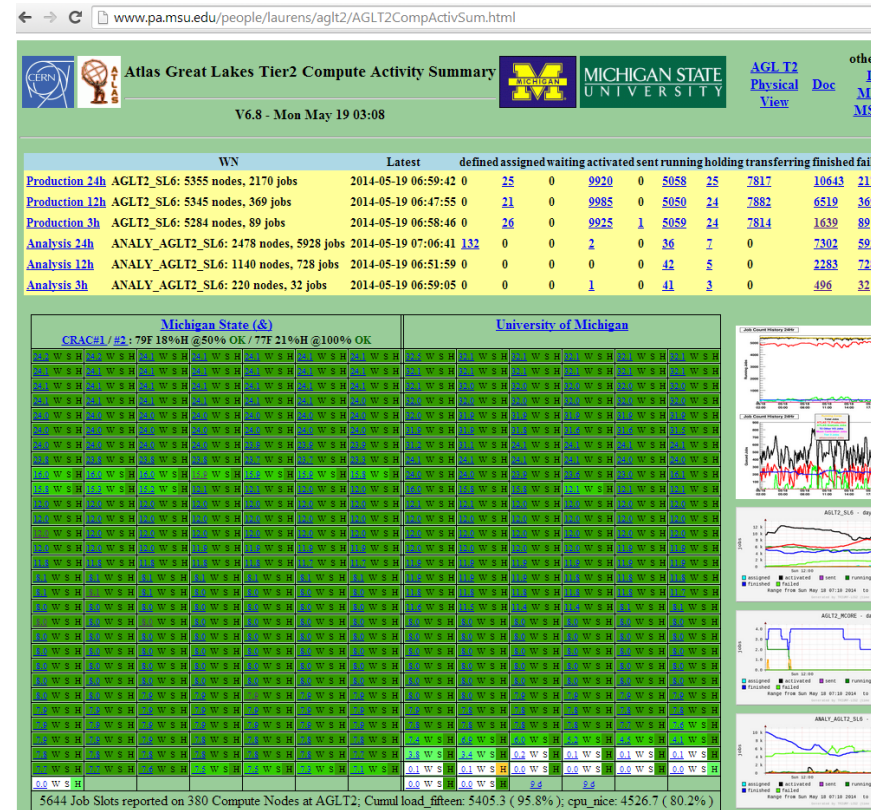
- ❏ I was notified in February that the University would be doing construction and renovation of the building which houses AGLT2 at the University of Michigan in 10 months.
 - ❏ Someone at the University determined that the space currently housing LS&A IT and AGLT2 would be required to host new HVAC equipment.
 - ❏ However they didn't check all the paperwork: We have a signed agreement from the University as part of the Tier-2 grant process to occupy that specific room for 5 more years.
- ❏ **We are discussing options but are in a strong position if there are disagreements about how to proceed.**
 - ❏ No single location meets our needs: equivalent power, space and cooling AND walk-able from physics (Undergrad access)
 - ❏ One option: split to two locations...under discussion
 - ❏ No costs to project. May involve some downtime(s)...

AGLT2 Monitoring

AGLT2 has a number of monitoring components in use

As shown in before we have:

- Customized “summary” page->
- **OMD (Open Monitoring Distribution)** at both UM/MSU
- **Ganglia**
- Central syslog'ing via ELK: Elasticsearch, Logstash, Kibana
- **SRMwatch** to track dCache SRM status
- **GLPI** to track tickets (with FusionInventory)



AGLT2 Provisioning/Config Mgmt

- AGLT2 uses a Cobbler server configuration managed by CFEngine and duplicated at both sites for building service nodes (excepting site-specific network/host info)
 - Created flexible default kickstart template with Cobbler's template language (Cheetah) to install a variety of "profiles" as selected when adding system to Cobbler (server, cluster-compute, desktop, etc).
 - Simple PXE based installation from network
 - Cobbler manages mirroring of OS and extra repositories
- Kickstart setup is kept minimal and most configuration done by CFEngine on first boot
- Dell machines get BIOS and Firmware updates in post-install using utils/packages from Dell yum repositories

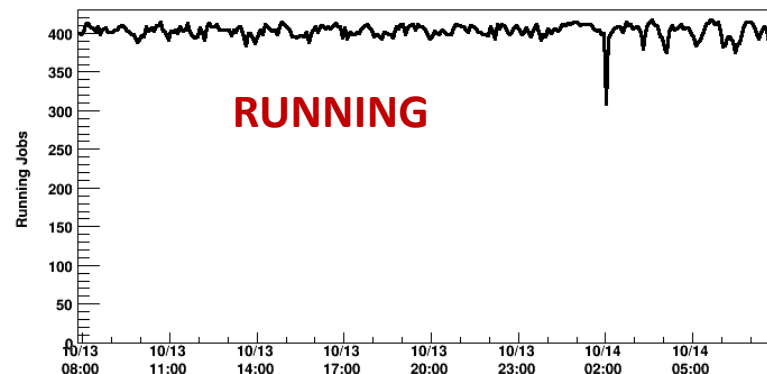
HTCondor CE at AGLT2

- Bob Ball worked for ~1 month at AGLT2 setup
 - Steep learning curve for newbies
 - Lots of non-apparent niceties in preparing job-router configuration
 - RSL no longer available for routing decisions
- Cannot change content of job route except during condor-ce restart
- However, CAN modify variables and place them in ClassAd variables set in the router
 - Used at AGLT2 to control MCore and LargeMemory slot access
 - In place on all gatekeepers
- See full details of our experience and setup at <https://www.aglt2.org/wiki/bin/view/AGLT2/CondorCE>

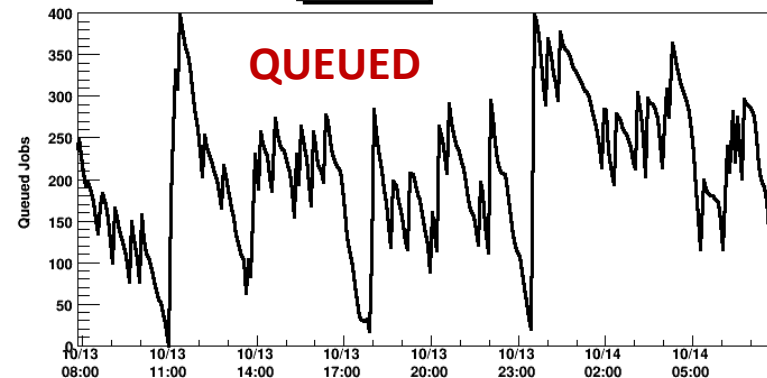
MCORE at AGLT2

- AGLT2 has supported MCore jobs for many months now
- Condor configured for two MCore job types
 - **Static** slots (10 total, 8 cores each)
 - **Dynamic** slots (420 of 8 cores each)
- Requirements statements added by the “condor_submit” script
 - Depends on count of queued MP8 jobs
- Result is instant access for a small number with gradual release of cores for more with time.

Job Count History 24Hr



Job Count History 24Hr



Full details at <https://www.aglt2.org/wiki/bin/view/AGLT2/MCoreSetup>

Large MEMory at AGLT2

- AGLT2 has supported LMEM jobs for the past few months now
- Condor configured for 6GB image allocation
- Machines dynamically configured to support LMEM slots
- Number per machine limited based upon hardware configuration
 - Utilizes an allocatable resource on each machine
- Requirements statements added by the HTCondor-CE Job Router
- **AGIS (ATLAS info system) must be defined with a floor in the memory**
 - Queue accepts jobs in the range of 4GB-6GB
 - Minimum is required or all jobs can (and will) use precious LMEM slots

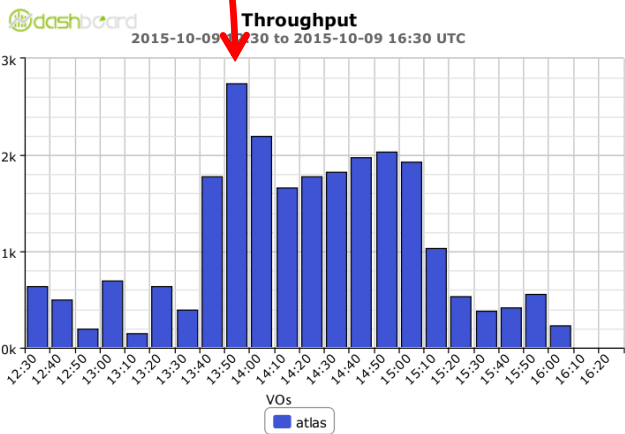
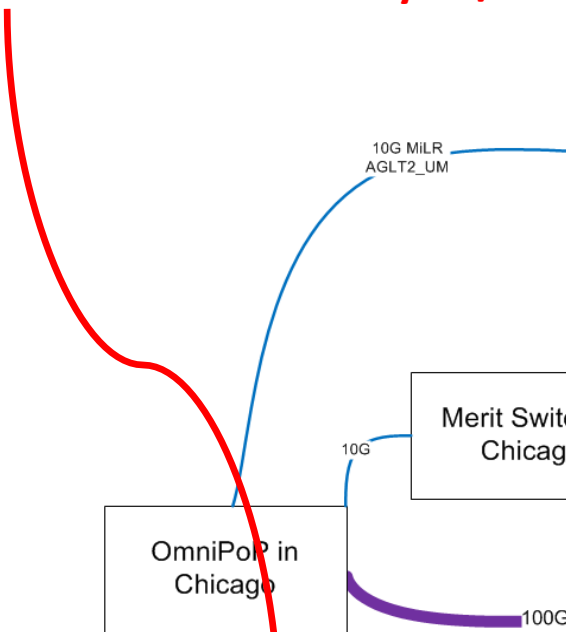
ATLAS Calibration

- There are 3 ATLAS calibration sites: **AGLT2** (Ann Arbor), **CNAF** (Rome), and **LMU** (Munich)
- This year all 3 sites switched over to replicating calibration job outputs using tool written at AGLT2 (Ben Meekhof)
 - Calib jobs write constants to Oracle database at each site which then must be put into central CERN db
 - Previously used Oracle Streams...placed constraints on database structure, large database changes, and had to be monitored closely
- Moving away from streams replication allowed us to **redo the database schema** with a more logical layout
 - Data was split into individual schemas for each year, now one unified schema with partitioned tables.
 - CERN and UM using new schema, CNAF considering, unsure at LMU
 - Database tool can handle either schema in source/destination

AGLT2 100G Network Details




100G to WAN works

Last Fall: "normal" FTS transfer hit 2.73 Gbytes/sec



Madalert

New project: high-level analysis on maddash meshes





-  Find problem patterns
-  Produces high level report
-  Useful when the grid is too big and “busy” to see problems

Started as a series of python scripts.

Being re-architected

-  Java web service for the analysis for future integration in MadDash
-  Client Javascript web UI
-  Python client for integration in OMD/check_mk and other tools

Rules for infrastructure problems

Pattern	Rule name	Description
	All down	All tests for all sites are unknown. This typically happens when the data store cannot be accessed.
	Site down	All tests for a particular site are unknown. This can happen if the site is down (check ping), the test service is down or the firewall is misconfigured.
	Site can't test	All tests initiated by a particular site are unknown. This can happen if the archive service is down, the scheduling service is doing, the configuration is faulty or the outgoing firewall rules are misconfigured.
		

Deployed: <http://maddash.aglt2.org/madalert.html> Source: <https://github.com/gcarcassi/madalert>

Future Plans

- Participating in SC16 (simple infrastructure for 100G+)
- Still exploring OpenStack as an option for our site. Testing Ceph for a back-end.
- New network components support Software Defined Networking (OpenFlow). We are experimenting with SDN in our Tier-2 and as part of LHCONE point-to-point testbed.
- Working on IPv6 dual-stack for all nodes in our Tier-2
 - Have IPv6 address block
 - Waiting for UM/MSU network engineers to complete routing (ESnet peering underway)

Summary

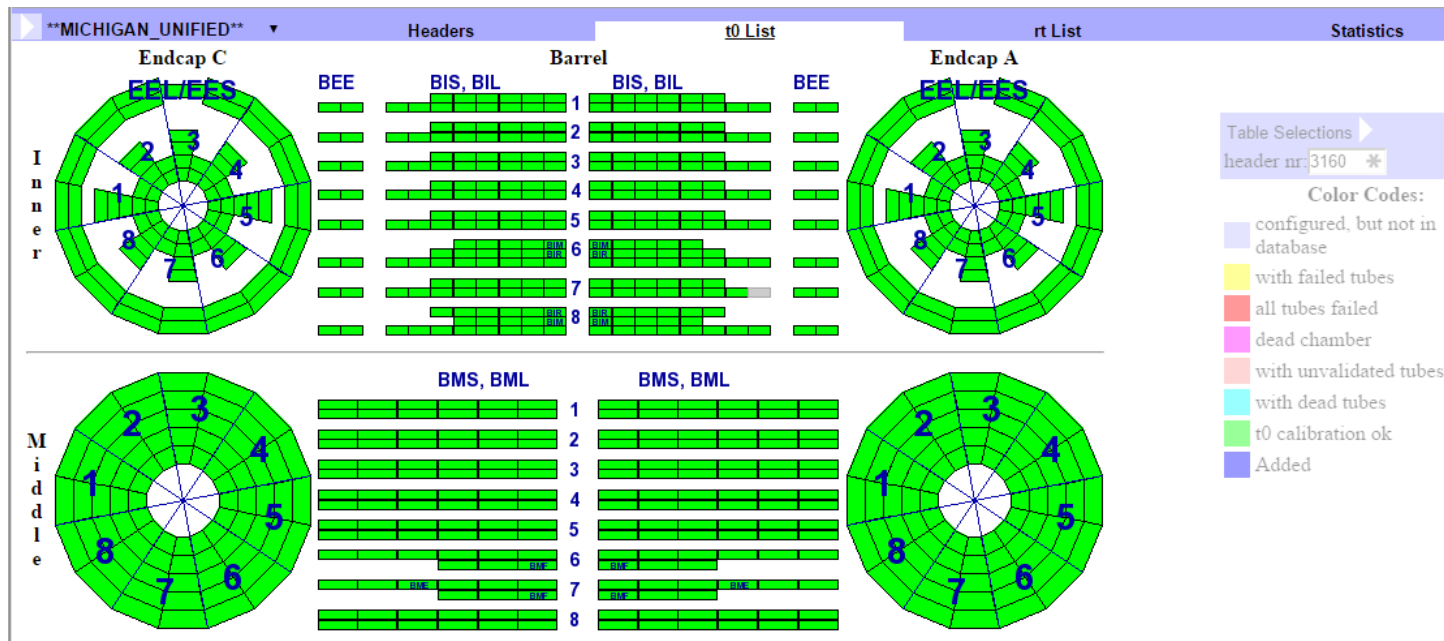
- Tier-2 services and tools are evolving. Site continues to expand and operations are smooth.
- **New large memory queue and multicore support enabled for AGLT2 using HT-Condor CE**
- **Muon calibration work completed to support Run-2 mode of operation. Working well**
- **Network 2x40G inter-site, 100G WAN, just works.**
- **MadAlert is a new project to analyze perfSONAR results and identify issues.**
- **FUTURE: OpenStack, IPv6, SDN**

Questions ?

Extra Slides

ATLAS Calibration

- With addition of **new chambers** to ATLAS the **Database Observer** used for visual overview of calibration had to be updated.
 - We took on maintenance of this project (originally written at Munich)
 - Did fixes to related package MuonFixedIdUnpack (also used by database tool)
 - SVN Web: https://svnweb.cern.ch/trac/msd/browser/calib_db_checks/trunk
 - All 3 sites + CERN viewable on our observer: <https://linat05.grid.umich.edu/calibdb>



Virtualization at AGLT2

- ❏ Most Tier-2 services run on VMware (vSphere 5.5u3)
- ❏ UM uses iSCSI storage backends
 - ❏ Dell MD3820i + MD1220, MD3600i + MD1200 and SUN NAS 7410
 - ❏ vSphere manages virtual disk allocation between units and RAID volumes based on various volume performance capabilities and VM demand
- ❏ MSU runs on DAS – Dell MD3200
- ❏ Site resiliency details
 - ❏ Multisite SSO operational between sites (SSO at either site manages both sites)
 - ❏ MSU is operating site-specific Tier-2 VMs (dcache doors, xrootd, cobbler) on vSphere
 - ❏ VMware Replication Appliance is used to perform daily replications of critical UM VMs to MSU's site. This is working well.

Madalert

Madalert Report

Mesh name: Latency tests between all WLCG hosts - Latency Tests Between WLCG Latency Hosts

Mesh location: <http://maddash.aglt2.org/maddash/grids/Latency tests between all WLCG hosts - Latency Tests Between WLCG Latency Hosts>

Infrastructure problems	
Site	Description
BelGrid-UCL_ingrid-ps01	Site can't test
CIEMAT-LCG2_perflat	Site can't test
DESY-HH_perfsonar-ps-01	Site mostly can't test
GRIF_perfsonar01	Site can't test
INFN-FRASCATI_psfrascati02	Site is down
IITP_psum02	Site can't be tested
KR-KISTI-GSDC-01_ps-gsdc01	Site mostly can't be tested
RO-14-TIM_sonar2	Site can't test
RU-SPbSU_alice14	Site can't test
Taiwan-LCG2_lhc-latency	Site can't test
UKI-SCOTGRID-DURHAM_perfmon	Site mostly can't test
USC-LCG2_perfsonar-ps-latency	Site can't test
UTA_SWT2_perfSONAR_latency	Site is down

Test failures	
Site	Description
BEgrid-ULB-VUB_ps01	Incoming tests failure (Loss rate is ≥ 0.01)
Caltech PerfSonar Latency	Incoming tests failure (Loss rate is ≥ 0.01)



Report for this grid →

Madalert

Report integrated in OMD/check_mk

Each site is a separate entry

Site count for green/yellow/red/orange is tracked over time

Check MK 1.2.4p5

Tactical Overview

Hosts: 273 Problems: 18 Unhandled: 18

Services: 4080 Problems: 1132 Unhandled: 1132

Quicksearch

localhost

Views

- ▶ Dashboards
- ▼ Hosts
 - All hosts
 - All hosts (Mini)
 - All hosts (tiled)
 - Favourite hosts
 - Host search
- ▶ Hostgroups
- ▶ Services
- ▶ Servicegroups
- ▶ Business Intelligence
- ▶ Problems
- ▶ Addons
- ▶ Other

Bookmarks

Add Bookmark

© Mathias Kettner

MAN-HEP_personar-it						
madalert_latency_UKI-SCOTGRID-DURHAM_perfmon	CRIT	CRIT - Site mostly can't test	3 hrs	38 sec		
madalert_latency_UKI-SCOTGRID-ECDF_gridpp-ps-lat	OK	OK - No match	3 hrs	38 sec		
madalert_latency_UKI-SCOTGRID-GLASGOW_ps002	CRIT	CRIT - Outgoing tests failure (Loss rate is >= 0.01)	2 hrs	38 sec		
madalert_latency_UKI-SOUTHGRID-BHAM-HEP_epgper	OK	OK - No match	3 hrs	38 sec		
madalert_latency_UKI-SOUTHGRID-CAM-HEP_serv04	OK	OK - No match	54 min	38 sec		
madalert_latency_UKI-SOUTHGRID-OX-HEP_t2ps-latency	OK	OK - No match	3 hrs	38 sec		
madalert_latency_UKI-SOUTHGRID-RALPP_heplnx129	CRIT	CRIT - Incoming tests failure (Loss rate is >= 0.01)	49 min	38 sec		
madalert_latency_UKI-SOUTHGRID-SUSX_grid-personar	OK	OK - No match	3 hrs	38 sec		
madalert_latency_US-AGLT2_MSU_LAT	CRIT	CRIT - Outgoing tests failure (Loss rate is >= 0.01)	96 min	38 sec		
madalert_latency_US-AGLT2_UM_LAT	OK	OK - No match	3 hrs	38 sec		
madalert_latency_US-FNAL_LT	OK	OK - No match	3 hrs	38 sec		
madalert_latency_US-MWT2_IU_LAT	OK	OK - No match	3 hrs	38 sec		
madalert_latency_US-MWT2_UC_LAT	OK	OK - No match	3 hrs	38 sec		
madalert_latency_US-MWT2_UIUC_LAT	OK	OK - No match	3 hrs	38 sec		
madalert_latency_US-Nebraska_LT	OK	OK - No match	3 hrs	38 sec		
madalert_latency_US-PNNL_LT	OK	OK - No match	3 hrs	38 sec		
madalert_latency_US-Purdue_LT	OK	OK - No match	3 hrs	38 sec		
madalert_latency_US-UCSD_LT	OK	OK - No match	3 hrs	38 sec		
madalert_latency_US-Vanderbilt_LT	CRIT	CRIT - Incoming tests failure (Loss rate is >= 0.01)	3 hrs	38 sec		
madalert_latency_US-Wisconsin_LT	OK	OK - No match	3 hrs	38 sec		
madalert_latency_USC-LCG2_perfsonar-ps-latency	CRIT	CRIT - Site can't test	3 hrs	38 sec		
madalert_latency_UTA_SWT2_perfSONAR_latency	CRIT	CRIT - Site is down	3 hrs	38 sec		
madalert_latency_wuppertalprod_personar	CRIT	CRIT - Outgoing tests failure (Loss rate is >= 0.01)	3 hrs	38 sec		
Memory used	OK	OK - 9.20 GB used (8.00 GB RAM + 1.11 GB SWAP + 0.09 GB Pagetables, this is 78.9% of 11.67 GB	2015-07-13 11:03:29	38 sec		

localhost: madalert_latency_US-AGLT2_MSU_LAT - count_0

current: 29.42 max: 60.00 avg: 47.15

localhost: madalert_latency_US-AGLT2_MSU_LAT - count_1

current: 42.40 max: 52.00 avg: 41.40

localhost: madalert_latency_US-AGLT2_MSU_LAT - count_2

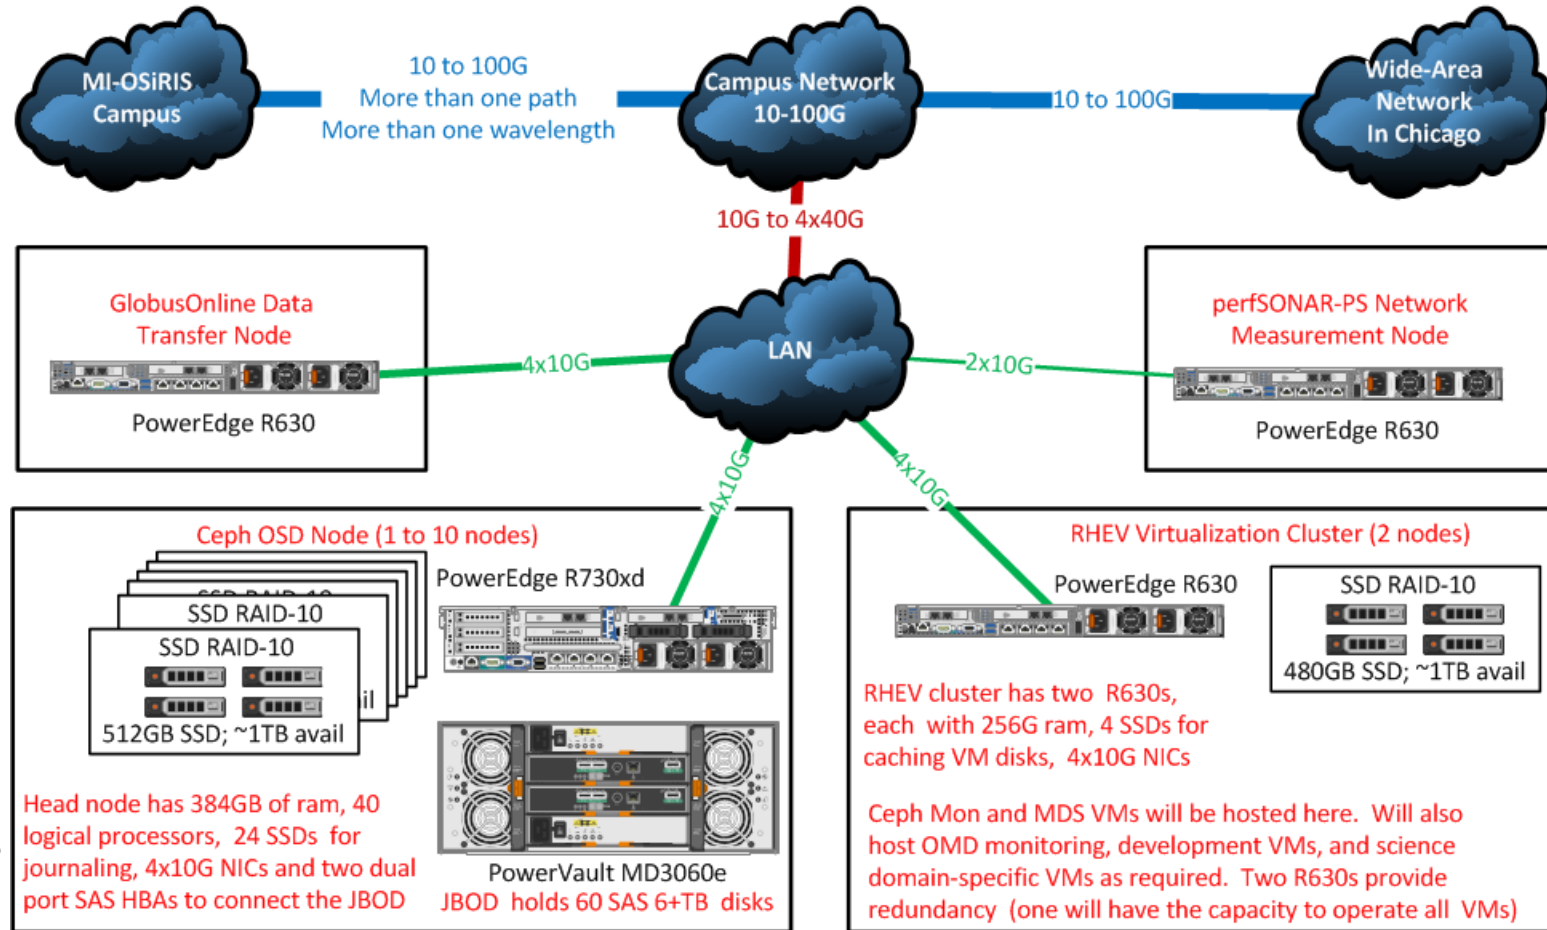
current: 96.42 max: 97.00 avg: 80.62

localhost: madalert_latency_US-AGLT2_MSU_LAT - count_3

current: 17.00 max: 17.00 avg: 17.00

Software-Defined Storage Research

MI-OSiRIS Data Infrastructure Building Block



See NSF announcement

http://www.nsf.gov/awardsearch/showAward?AWD_ID=1541335

NSF proposal
OSiRIS **funded**
involving UM,
MSU, WSU (and
our Tier-2)

Exploring
Ceph+SDN for
future software-
defined storage

Goal is
centralized
storage that
supports in place
access from CPUs
across multiple
institutions