



AGLT2 Site Report










Shawn McKee

University of Michigan

HEPiX Spring 2014



Outline

-  Site Summary and Status
-  Monitoring
-  Provisioning with Cobbler
-  Configuration with CFEngine 3
-  MCore on AGLT2
-  AFS on Linux ZFS
-  Virtualization Status
-  Networking Upgrade
-  Conclusion

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
 - ❏ 5331 single core job slots
 - ❏ 10 dedicated eight-core (multi-core) job slots
 - ❏ 233 Tier-3 job slots usable by Tier-2
 - ❏ Average 9.03 HS06/slot
 - ❏ 3.5 Petabytes of storage
- ❏ Most Tier-2 services virtualized in VMware
- ❏ 10Gb redundant ring for inter-site connectivity, lots of 10Gb internal ports and 12 x 40Gb ports
- ❏ High capacity storage systems have 2 x 10Gb bonded links
- ❏ 40Gb link between Tier-2 and Tier-3 physical locations

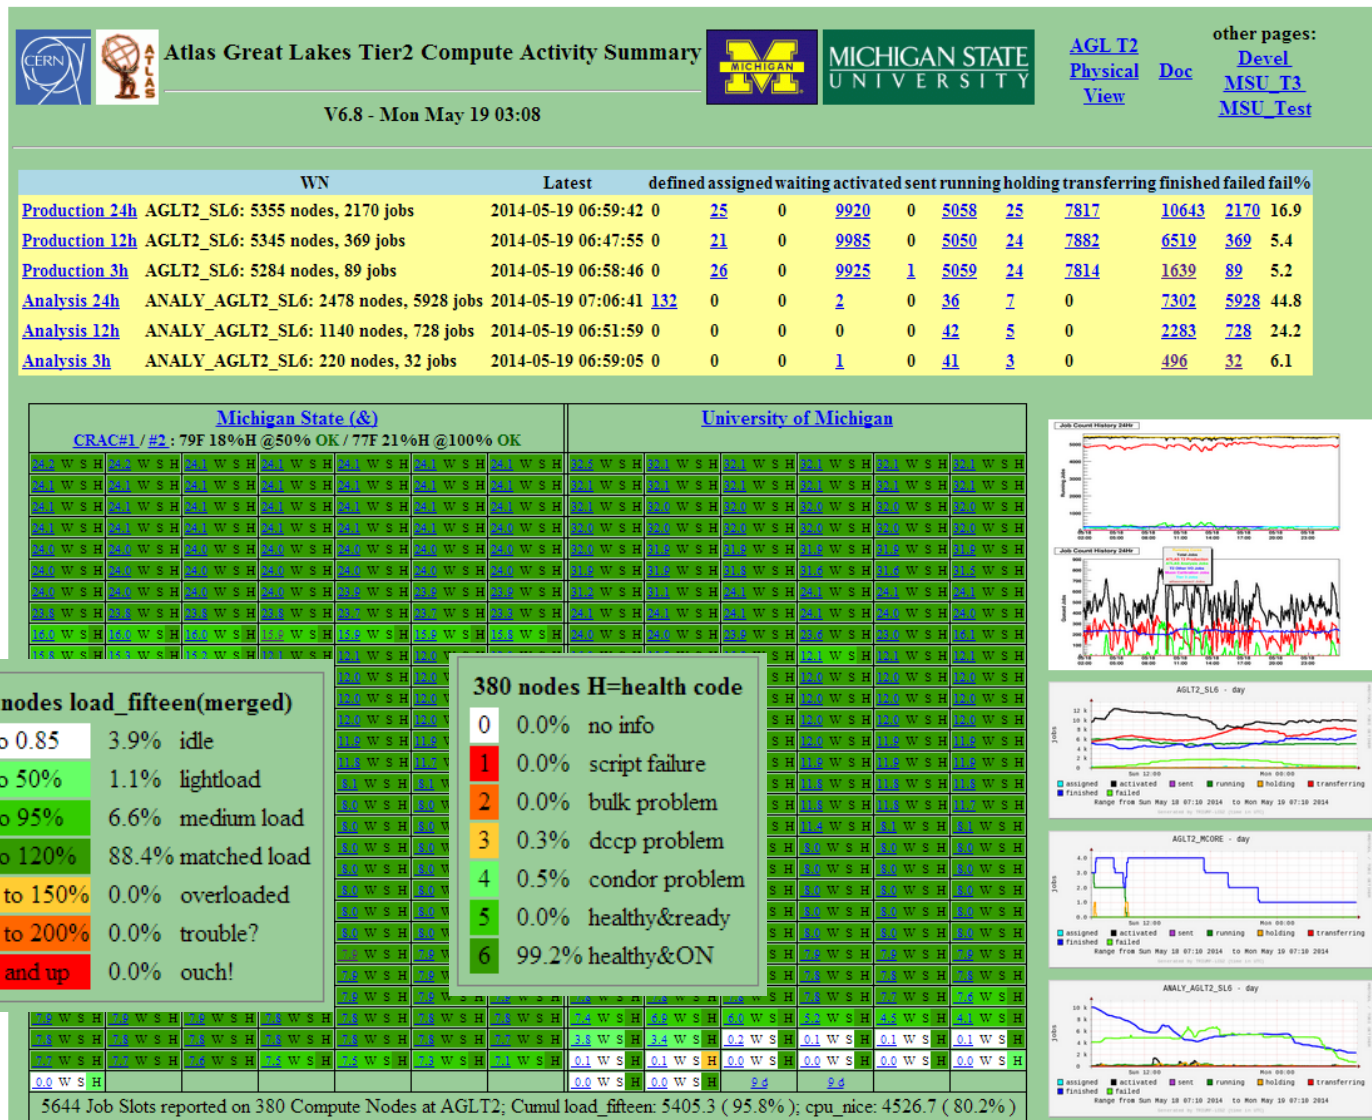
Custom Site Summary Monitoring

← → ↻ www.pa.msu.edu/people/laurens/aglt2/AGLT2CompActivSum.html

AGLT2 uses a custom page to summarize status shown here

Each node is colored in 4 bins (dark-green is "good")

Clicking drills-down to Ganglia or custom graphs



New Monitoring using OMD

AGLT2 has started using a new monitoring system to track the status of various services: OMD (Open Monitoring Distribution)

OMD (see <http://omdistro.org/start>) bundles Nagios variants (Nagios, Shinken, Icinga) and common tools (Check_MK, Nagvis, Thruk, PNP4Nagios) into a single distribution preconfigured to work with one another.

Easy to deploy (one RPM)

Straightforward to configure

Automates finding services and setting up graphs

Next few pages show an example from AGLT2's UM site....

OMD Check_MK Hostgroup Example

Check_MK 1.2.4 Hostgroups (Summary) 20 rows omdadmin (admin) 01:45

Tactical Overview

Hosts	Problems	Unhandled
359	4	4

Services	Problems	Unhandled
10083	1502	1502

Quicksearch

Bookmarks

Views

- Dashboards
 - Host & Services Problems
 - Main Overview
 - Network Topology
- Hosts
 - All hosts
 - All hosts (Mini)
 - All hosts (tiled)
 - Favourite hosts
 - Host search
- Hostgroups
 - Hostgroups
 - Hostgroups (Grid)
 - Hostgroups (Summary)
- Services
 - All services
 - Favorite services
 - Recently changed services
 - Serv. by host groups
 - Service search
- Servicegroups
- Business Intelligence
- Problems
 - Alert Statistics
 - Host problems
 - Pending Services

Hostgroups Table

Name	Alias	Up	Dw	Un	Pd	O	W	C	U	P
AFSNodes	AFSSystems	7	0	0	0	195	0	0	0	0
DYNES	DYNES	2	0	0	0	157	0	1	0	0
DellOMSA	DellOMSA	33	2	0	0	3919	5	40	1450	0
EnvironmentalMonitoring	TempHumidityNodes	11	0	0	0	158	0	0	0	0
Gatekeepers	GridPortals	3	0	0	0	99	1	0	0	0
InteractiveNodes	LoginNodes	4	0	0	0	291	0	0	0	0
LinuxHosts	LinuxHosts	49	2	0	0	4333	6	41	1450	0
LustreNodes	LustreSystems	11	0	0	0	949	1	15	916	0
MD3xxx	DelIMD3xxxNodes	4	0	0	0	44	0	0	0	0
NFSNodes	NFSNodes	5	0	0	0	241	0	3	69	0
PdUs	PdUs	20	0	0	0	516	0	0	0	0
RemoteAccess	RACs	191	1	0	0	876	0	2	0	0
Squids	CachingNodes	3	0	0	0	99	0	0	0	0
SwitchesRouters	NetDevices	20	1	0	0	1206	1	2	0	0
UMHosts	UMHosts	345	2	0	0	8237	9	44	1448	0
UPSeS	UPSeS	7	0	0	0	62	0	0	0	0
VMWareHosts	ESXi	3	0	0	0	141	0	0	0	0
VirtualMachines	VMwareVMs	50	0	0	0	1404	3	3	0	0
WorkerNodes	WorkerNodes	0	2	0	0	93	0	3	2	0
dCacheNodes	dCacheRelated	18	0	0	0	2350	4	19	463	0

refresh: 30 secs

OMD/Check_MK Host Details

omd.aglt2.org/atlas/check_mk/index.py?start_url=%2Fatlas%2Fcheck_mk%2Fview.py%3Fview_name%3Dhost%26host%3Dhead01%26site%3Ddomr

Check_MK 1.2.4

Services of Host head01, UM

95 rows omdadmin (admin) 01:47

Host status | Graphs | Host notifications | Host downtimes | Edit View

Availability

head01

State	Service	Icons	Status detail	Age	Checked	Perf-O-Meter
OK	Check_MK		OK - Agent version 1.2.4b6, execution time 1.5 sec	2014-05-09 14:52:02	47 sec	1.5s
OK	Check_MK inventory		OK - no unchecked services found	2014-05-09 14:52:47	29 sec	
OK	CPU load		OK - 15min load 0.49 at 16 CPUs	2014-02-25 12:51:52	46 sec	0.5
OK	CPU utilization		OK - user: 4.2%, system: 0.8%, wait: 0.0%	2014-02-25 12:51:52	46 sec	5%
OK	Dell OMSA		OK - System: 'PowerEdge R610 II', SN: '48 GB ram (6 dimms), 3 logical drives, 6 physical drives', BIOS='6.07 08/18/2011', iDRAC6='1.96.00 (Build 1)', Ctrl 0 [PERC H700 Integrated]: Fw='12.10.6-0001', Dr='06.700.06.00-rh1', End 0:0:0 [Backplane]: Fw='1.07', End 0:1:0 [Backplane]: Fw='1.07', OpenManage Server Administrator (OMSA) version: '7.4.0'	2014-05-09 14:52:44	32 sec	
OK	Disk IO SUMMARY		OK - 11.67kB/sec read, 233.47kB/sec write, IOs: 27.47/sec	2014-02-25 12:52:52	46 sec	0.01M/s 0.23M/s
OK	DMI Sysinfo		OK - Manufacturer: Dell Inc., Product-Name: PowerEdge R610, Version: Not Specified, S/N: 85V3KN1	2014-02-25 12:51:52	46 sec	
OK	ESM Log		OK - EMS Log is less than 80% full	2014-04-16 03:48:42	46 sec	
OK	fs_/_		OK - 65.1% used (323.94 of 497.2 GB), (levels at 90.00/96.00%), trend: +205.24MB / 24 hours	2014-02-25 12:51:53	46 sec	65.15%
OK	fs_/_boot		OK - 49.4% used (0.09 of 0.2 GB), (levels at 90.00/96.00%), trend: +0.00B / 24 hours	2014-02-25 12:51:53	46 sec	49.37%
OK	fs_/_tmp		OK - 6.2% used (0.91 of 14.8 GB), (levels at 90.00/96.00%), trend: -52.08B / 24 hours	2014-02-25 12:51:53	46 sec	6.18%
OK	fs_/_var		OK - 33.2% used (9.79 of 29.5 GB), (levels at 90.00/96.00%), trend: -2.58MB / 24 hours	2014-02-25 12:51:53	46 sec	33.17%
OK	fs_/_var/cache/sonata		OK - 5.9% used (0.23 of 3.9 GB), (levels at 90.00/96.00%), trend: +78.14kB / 24	2014-02-25 12:51:53	46 sec	5.90%

Host services can be auto-discovered via **check_mk_agent** or snmp
Corresponding graphs auto-created. CMK supports more than 500 services ('cmk -L')

Check MK “hover” example

← → ↻ omd.aglt2.org/atlas/check_mk/index.py?start_url=%2Fatlas%2Fcheck_mk%2Fview.py%3Fview_name%3Dhost%26host%3Dhead01%26site%3Domi ☆ ⚙️ ❤️ 📌 ✖️

Check **MK** 1.2.4 Services of Host head01, UM 95 rows omdadmin (admin) 01:47 **MK**

Availability 2 30s Host status Graphs Host notifications Host downtimes Edit View

head01

State	Service	Icons	Status detail	Age	Checked	Perf-O-Meter
OK	Check_MK		OK - Agent version 1.2.4b6, execution time 1.6 sec	2014-05-09 14:52:02	17 sec	1.6s
OK	Check_MK inventory		OK - no unchecked services found	2014-05-09 14:52:47	59 sec	
OK	CPU load		OK - 15min load 0.68 at 16 CPUs	2014-02-25 12:51:52	16 sec	2.3
OK	CPU utilization		OK - user: 11.6%, system: 15.1%, wait: 0.0%	2014-02-25 12:51:52	16 sec	26%
OK	Dell OMSA		OK - System: 'PowerEdge R610 II', SN: '48 GB ram (6 dimms), 3 logical drives, 6 physical drives', BIOS='6.0.7 08/18/2011', iDRAC6='1.96.00 (Build 1)', Ctrl 0 [PERC H700 Integrated]: Fw='12.10.6-0001', Dr='06.700.06.00-1', Encl 0:1:0 [Backplane]; rator (OMSA) version: '7.4.0'	2014-05-09 14:52:44	2 sec	
OK	Disk IO SUMMARY		115.45/sec	2014-02-25 12:52:52	16 sec	0.08M/s 0.73M/s
OK	DMI Sysinfo		PowerEdge R610, Version: Not	2014-02-25 12:51:52	16 sec	
OK	ESM Log		OK - EMS LOG is less than 80% full	2014-04-16 03:48:42	16 sec	
OK	fs_/_		90.00/96.00%, trend: +0.00B / 24	2014-02-25 12:51:53	16 sec	65.15%
OK	fs_/boot		90.00/96.00%, trend: -52.06B / 24	2014-02-25 12:51:53	16 sec	49.37%
OK	fs_/tmp		90.00/96.00%, trend: -2.54MB / 24	2014-02-25 12:51:53	16 sec	6.18%
OK	fs_/var		96.00/96.00%, trend: +78.10kB / 24	2014-02-25 12:51:53	16 sec	33.17%
OK	fs_/var/cache/openafs		90.00/96.00%, trend: +1.00MB / 24	2014-02-25 12:51:53	16 sec	5.90%
OK	fs_/var/lib/pgsql		90.00/96.00%, trend: +1.00MB / 24	2014-03-11 19:52:50	16 sec	64.86%
OK	Interface em1		13.94kB/s, out: 375.54kB/s	2014-02-25 12:51:52	16 sec	0.0% 0.3%
OK	Interface em2		126.22kB/s, out: 774.38kB/s	2014-02-25 12:51:52	16 sec	0.7% 0.6%
OK	IPMI Sensor			2014-02-25 12:51:52	16 sec	

Hovering over graph icon gives preview, clicking drills down to graph page.

AGLT2 Replaced PHP-Syslog-NG

AGLT2 had been using a central syslog-ng host and making the data searchable via php-syslog-ng.

As we migrated to SL6 from SL5 on our servers we felt it was time to update to something new and better supported.

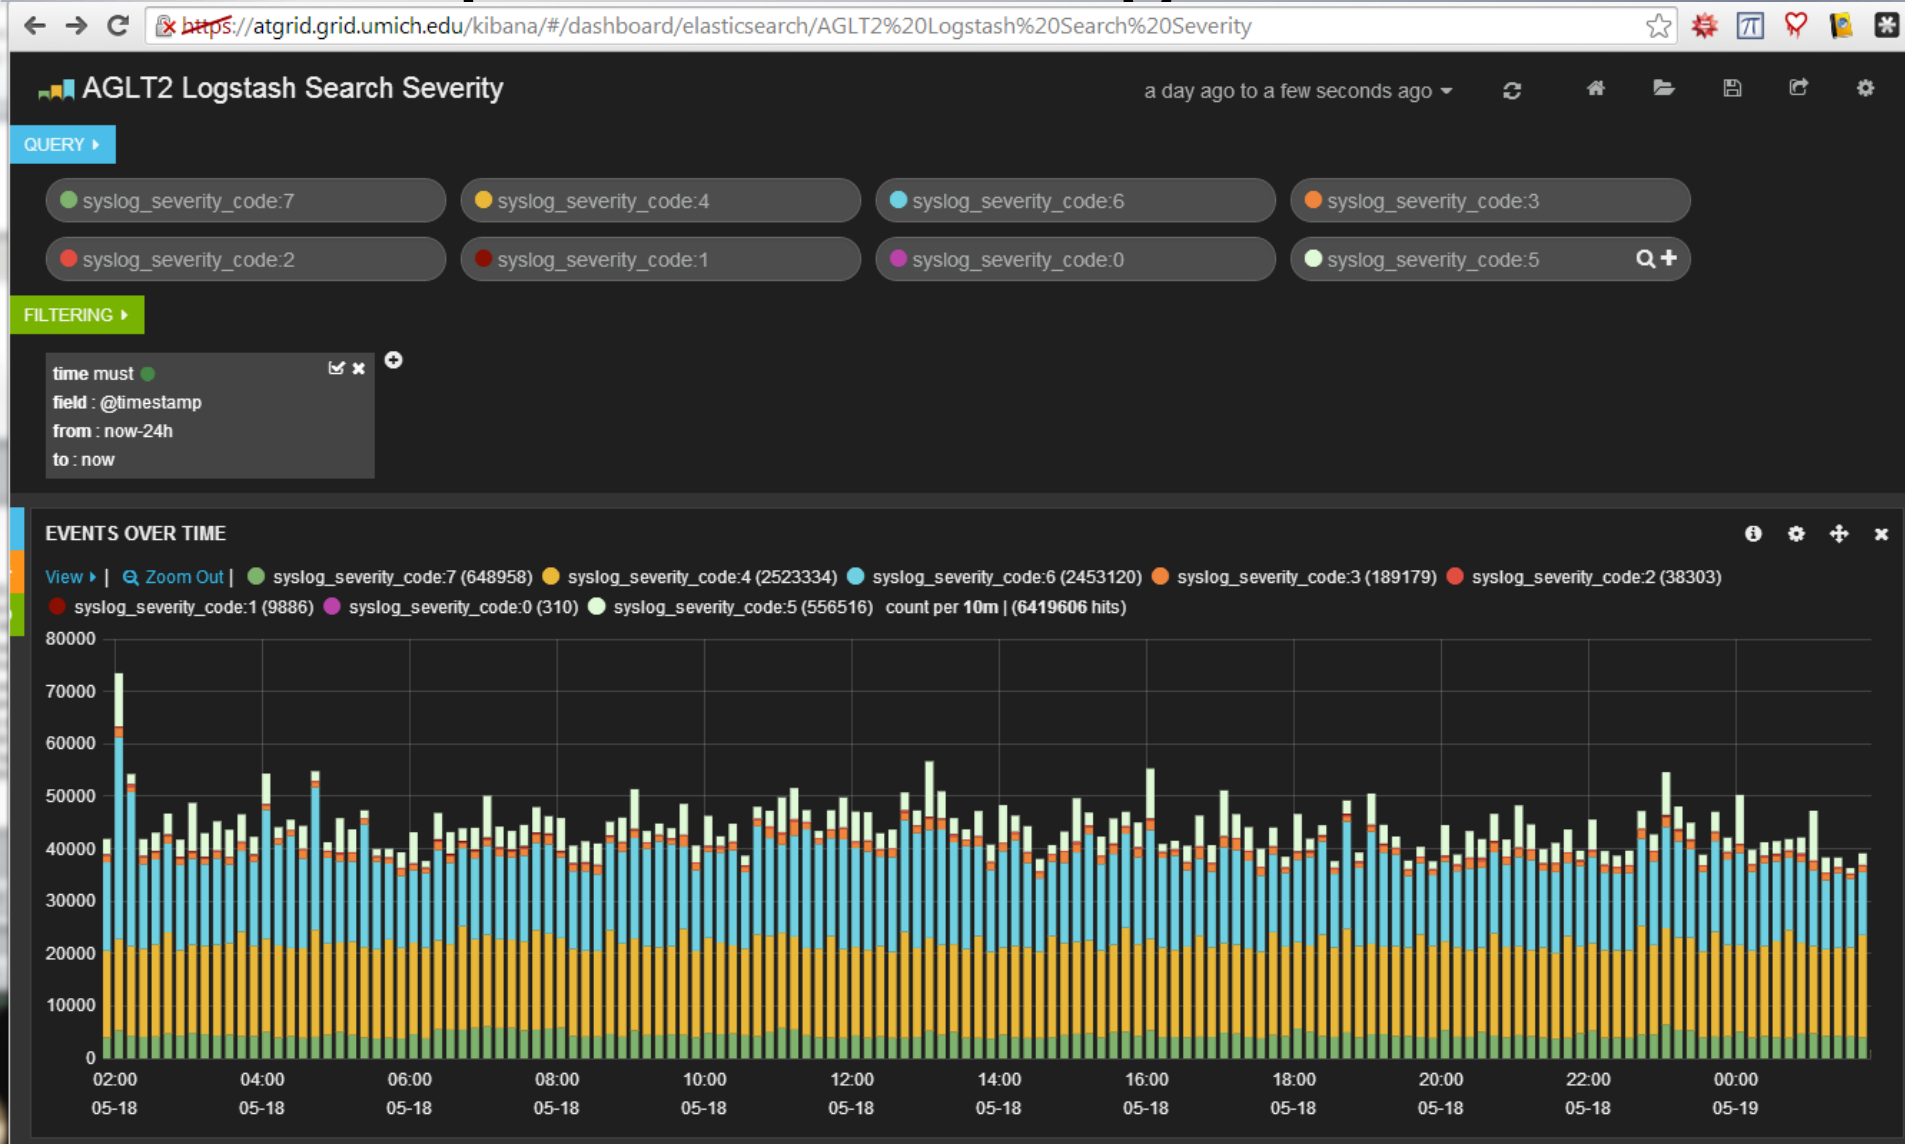
At the last HEPiX we saw sites using Logstash, Elasticsearch, Kibana and decided to try it out.

Initially had some issues related to installing a release-candidate and not having proper cleanup in place.

curl -XDELETE http://localhost:9200/<path_to_shards>

This runs over selected locations daily, deleting files from 1 week before

Example Kibana Page for AGLT2

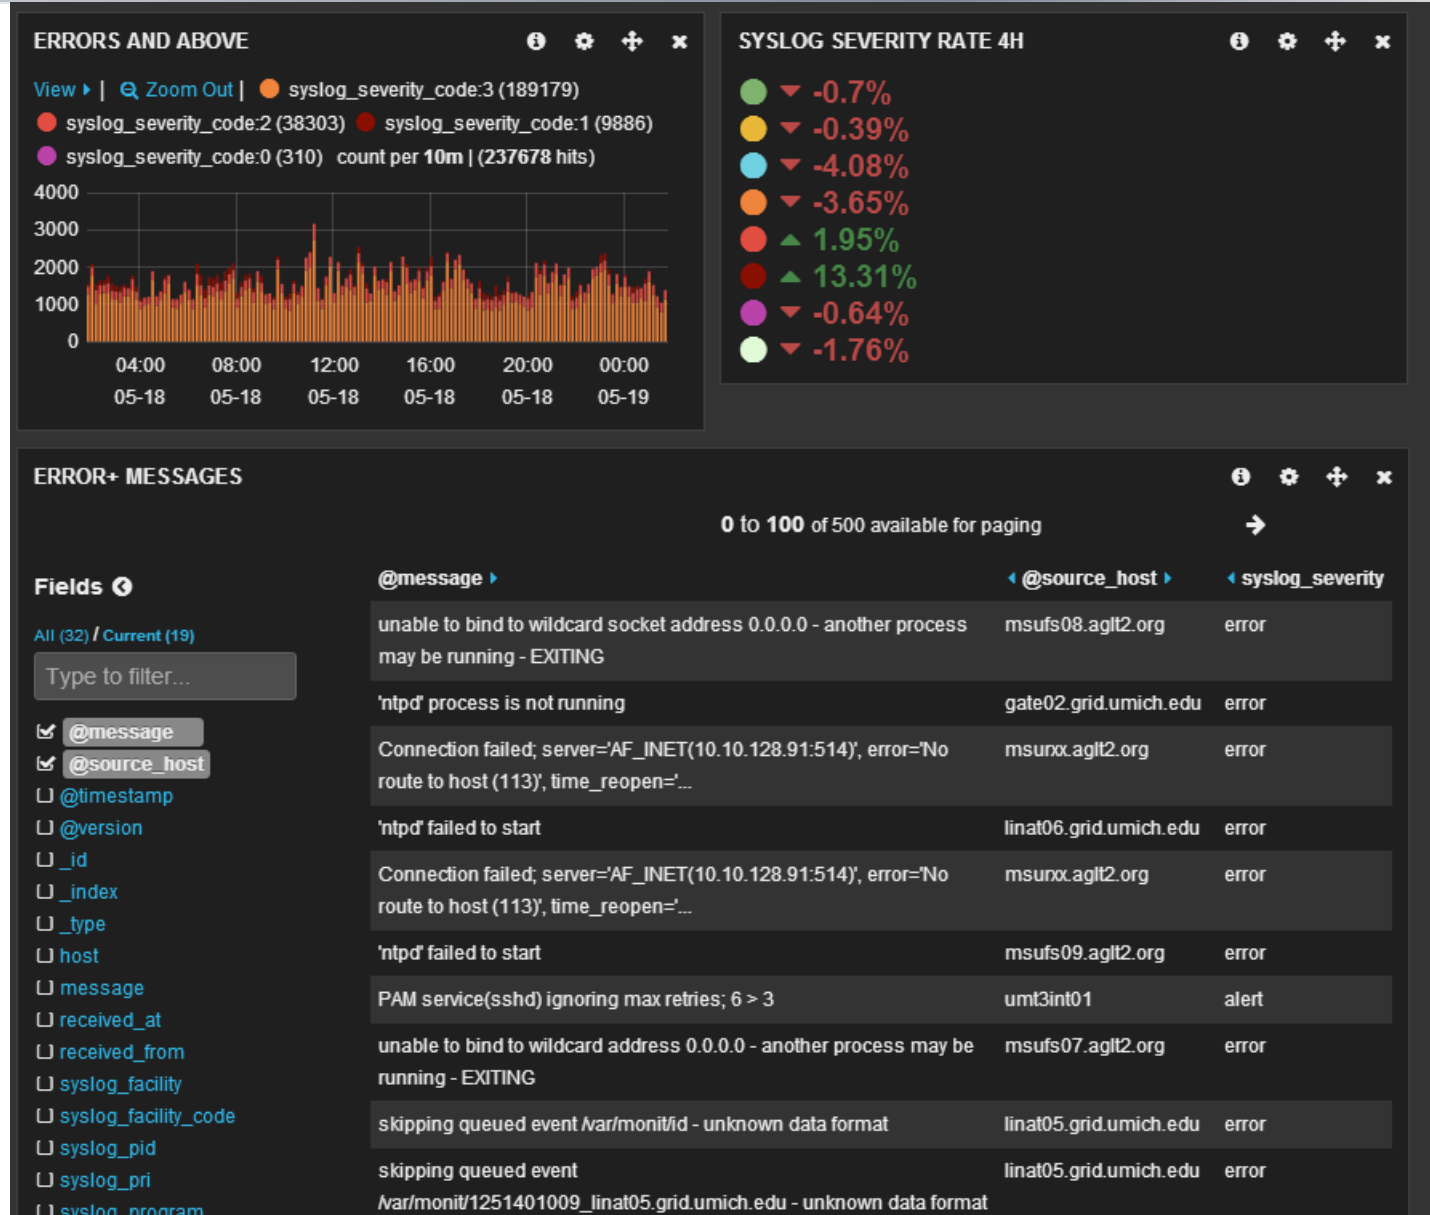


Highlight ERROR and above

Predefined views can be configured to focus on specifics.

This example shows all ERROR and above criticality events being recorded in our central syslog server.

All our switches, worker-nodes, servers and infrastructure components that can syslog send to our server



AGLT2 Provisioning

- 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 handles (with included post-install scripts) creating bonded NIC configurations – used to deal with those manually
 - 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
- Longer term we plan to replace Rocks with Cobbler for all compute nodes

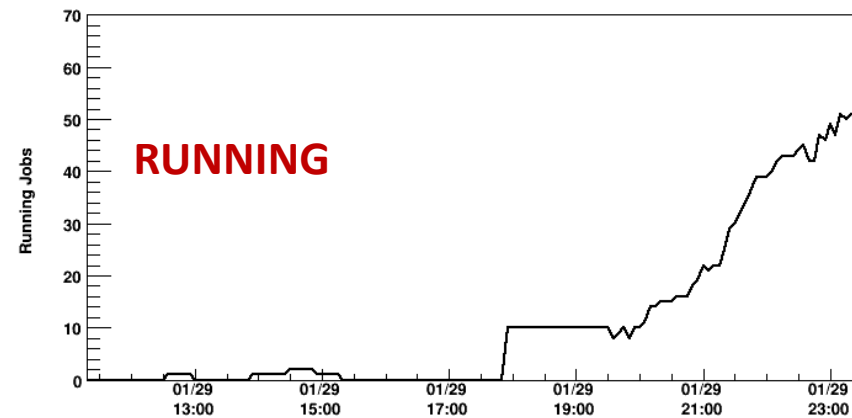
CFEngine 3 at AGLT2

- All systems are configured by CFEngine 3. MSU and UM sites have identically configured policy servers and policy.
 - Any client can contact either policy host for updates
- Developed workflow which easily allows any person to work on their own SVN branch and export it for testing from policy servers.
- CFEngine policy updates some system firmware in background – in particular RAC and Dell Lifecycle Controller which can cause other updates to fail validation if outdated.
 - These components require no manual intervention/reboot
- Once upon a time (3.0.x) there were bugs in iteration with array datatypes that stood in the way of many efficient approaches to policy code. Those are resolved in > 3.1 (we use current release 3.5.3).

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 (up to 230 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 12 Hours



Job Count History 12 Hours

MCORE-8 Jobs



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

Experience with AFS on ZFS

- 🖥️ At AGLT2 we run our own AFS cell `atlas.umich.edu`
- 🖥️ Once we invested in VMware we migrated our AFS services there
 - 🖥️ The three file servers at UM were virtualized and the `/vicepX` partitions were moved on iSCSI locations (still using `ext3`)
- 🖥️ Once ZFS on Linux was available we migrated to it for our AFS `/vicepX` storage: compression, snapshots and data integrity were primary motivations
 - 🖥️ Compression provides factor of **1.4** increase in space depending on partition
 - 🖥️ `'zfs get "all" zfs/vicepg | grep compressratio'`
- 🖥️ One issue: snapshots “steal” filesystem space! Have had “full” partitions a few times requiring us to manually fix. Just need to tune snapshots and balance space
 - 🖥️ Check space used via `'zfs list -t snapshot'`
- 🖥️ **Have recently hit an intermittent issue with ‘vos release’ in this setup**
 - 🖥️ OpenAFS 1.6.5.1-148 and ZFS 0.6.2 on SL 6.5
 - 🖥️ AFS is using DAFS (demand-attached file server) mode
 - 🖥️ When ‘vos release’ fails a forced salvage “fixes” things
 - 🖥️ Need to better document the problem and report it on the OpenAFS list
 - 🖥️ New 1.6.6+ AFS might help...

Virtualization at AGLT2

- ☐ Most Tier-2 services on VMware (now vSphere 5.5)
- ☐ UM uses iSCSI storage backends
 - ☐ Dell MD3600i, MD3000i 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
- ☐ Working on site resiliency details
 - ☐ vSphere and storage operational at MSU
 - ☐ Multisite SSO operational between sites (SSO at either site manages both sites)
 - ☐ Have enabled vSphere Replication for critical servers
 - ☐ Working on details of overall site resilient configuration, how to move services and inform upstream (or do transparently)

vSphere Web Client

SSO can be made using either site's local vCenter web interface

The screenshot displays the VMware vSphere Web Client interface. The browser address bar shows the URL: <https://umvmgt.physics.lsa.umich.edu:9443/vsphere-client/#extensionId=vsphere.core.folder.manage.scheduledOpsView;context=com.vmware.core>. The interface includes a left-hand navigation pane with a tree view showing the hierarchy: vCenter > msu-vcenter > MSU-WINSVC2 > umvmgt.physics.lsa.umich.edu > AGLT2 > AGLT2_UMCluster > umvm01.aglt2.org, umvm02.aglt2.org, umvm03.aglt2.org, and (85) Virtual Machines. The main content area is titled 'msu-vcenter' and has tabs for 'Getting Started', 'Summary', 'Monitor', 'Manage', and 'Related Objects'. The 'Manage' tab is active, and within it, the 'Scheduled Tasks' sub-tab is selected. A message states: 'To create a scheduled task, select an action from the Schedule New Task drop-down menu from below. You can also navigate to an object in the Inventory Lists, click the Actions menu, and press Ctrl. The clock icon that appears when you press CTRL indicates the actions that you can schedule on the object, such as Create Snapshot, or Add Host. Select an action and configure the scheduling options.' Below this message is a table with the following columns: Task, Schedule, Last Run, Last Run Result, and Next Run. The table is currently empty, displaying 'This list is empty.' in the 'Last Run' column. At the bottom of the table, it says '0 items'. On the right side of the interface, there are two panels: 'Recent Tasks' and 'Work In Progress'. The 'Recent Tasks' panel shows a list of tasks, including 'Check new notifications' and 'umvmgt.physics.lsa.umich.edu', with status indicators for 'All', 'Running', and 'Failed'. The 'Work In Progress' panel shows tasks with warning icons, including 'umvmgt.physics.... (4)', 'cache4 - Migrate', 'AGLT2_UMClust...', and 'umvmgt.physics.... (3)'. At the bottom of the interface, it says 'No items selected'.

Working well for most VM / Datastore management tasks.

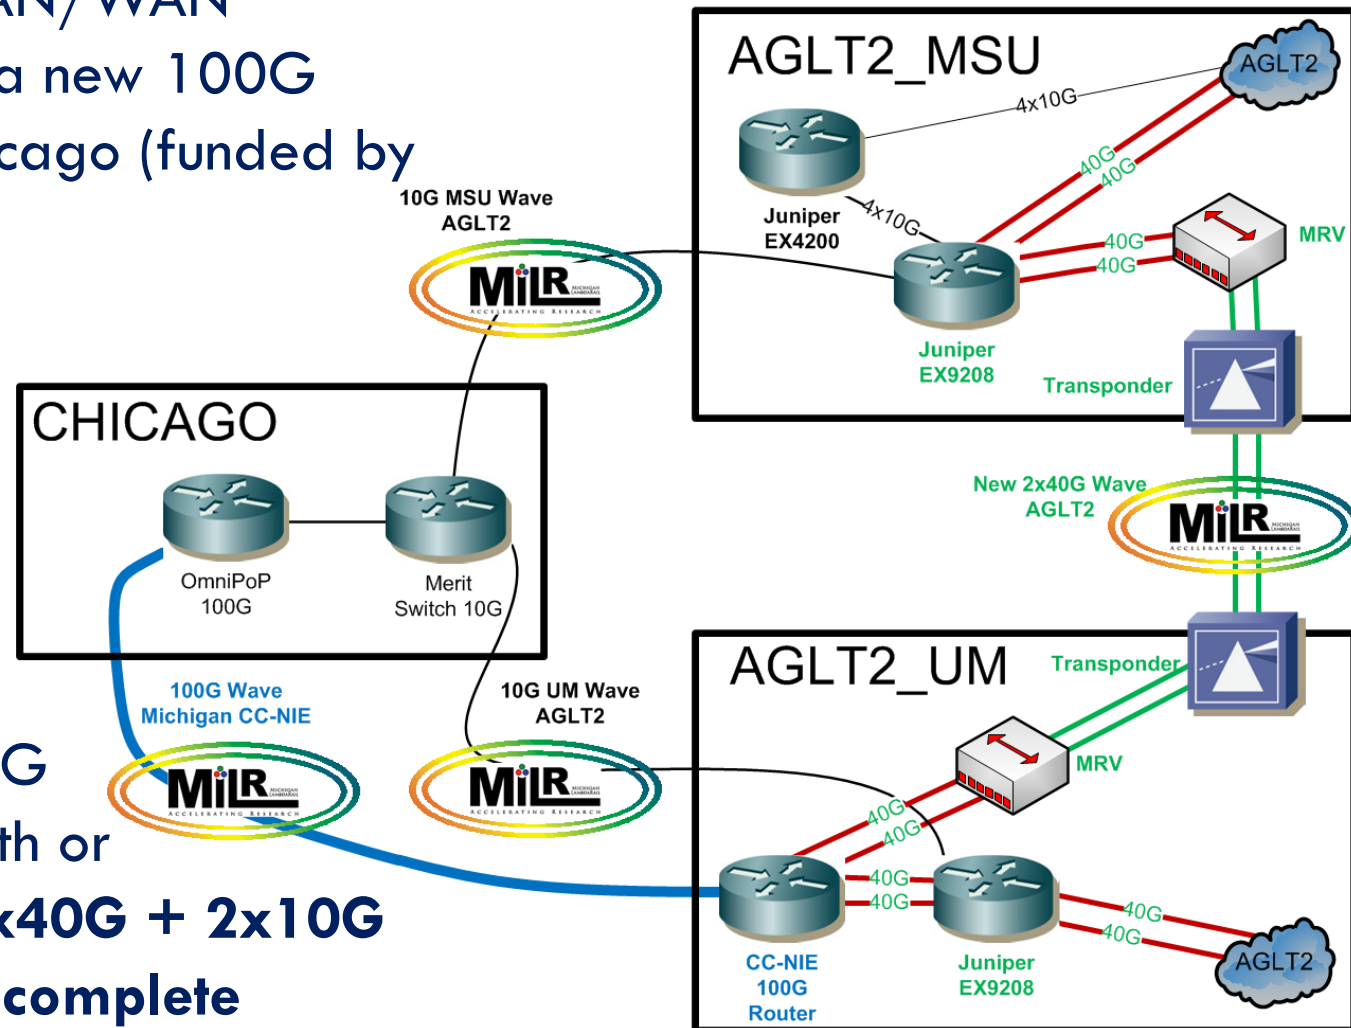
Site Resiliency Using VMware

- ❏ Our goal is to have MSU capable of bringing up Tier-2 service VMs within 1 day of loss of UM site
 - ❏ 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
- ❏ VMware Replication Appliance does not currently support automatic failover/failback
 - ❏ To failover, VMs will need to be started and network and VM configuration will need to be modified manually.
 - ❏ If outage is long, VMs may need to failback to UM, which also requires manual intervention.

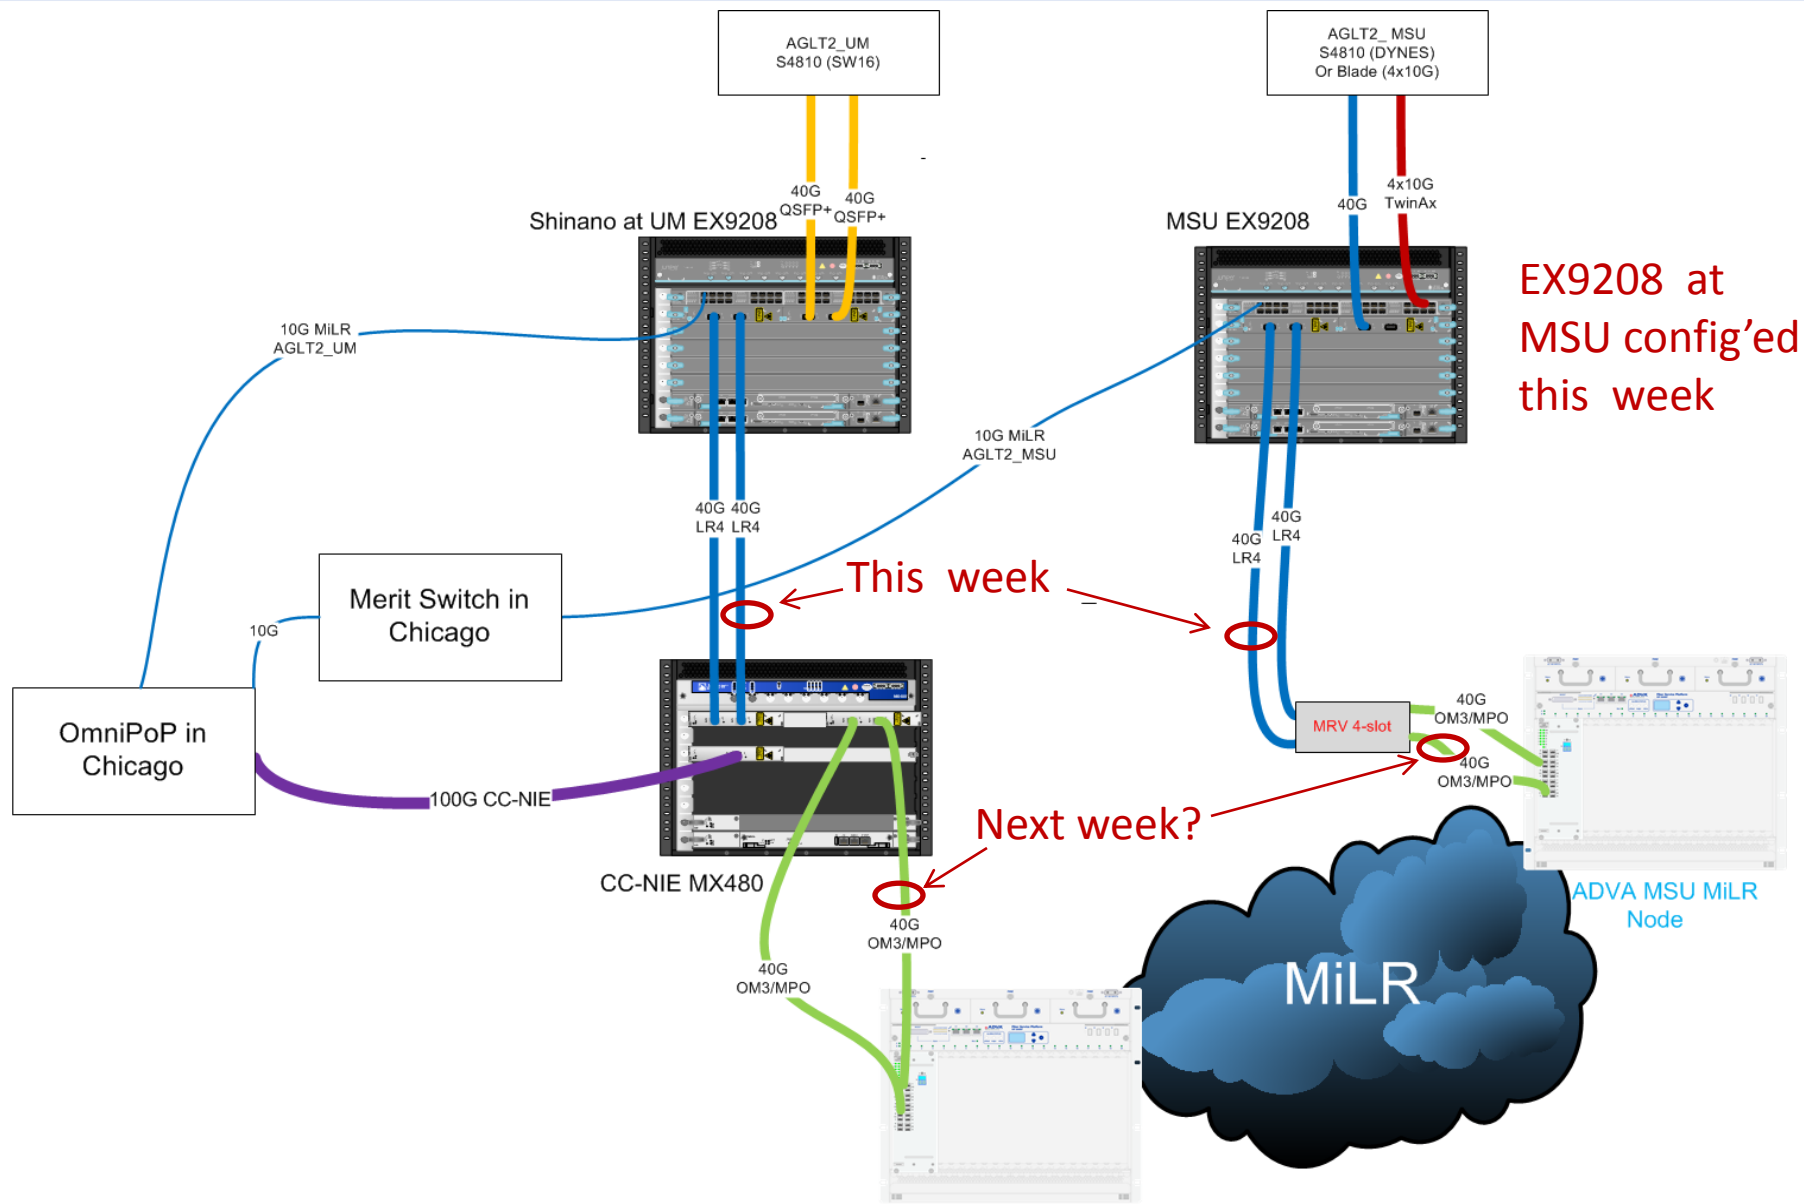
AGLT2 Network Upgrade

AGLT2 Network Upgrade Diagram

- AGLT2 is in the process of upgrading our LAN/WAN
- U Michigan has a new 100G connection to Chicago (funded by NSF CC-NIE)
- Tier-2 used project funds to purchase Juniper EX9208 routers at UM/MSU
- Deploying 2x40G over the next month or two. **Will have 2x40G + 2x10G to the WAN once complete**



AGLT2 Net-upgrade Details



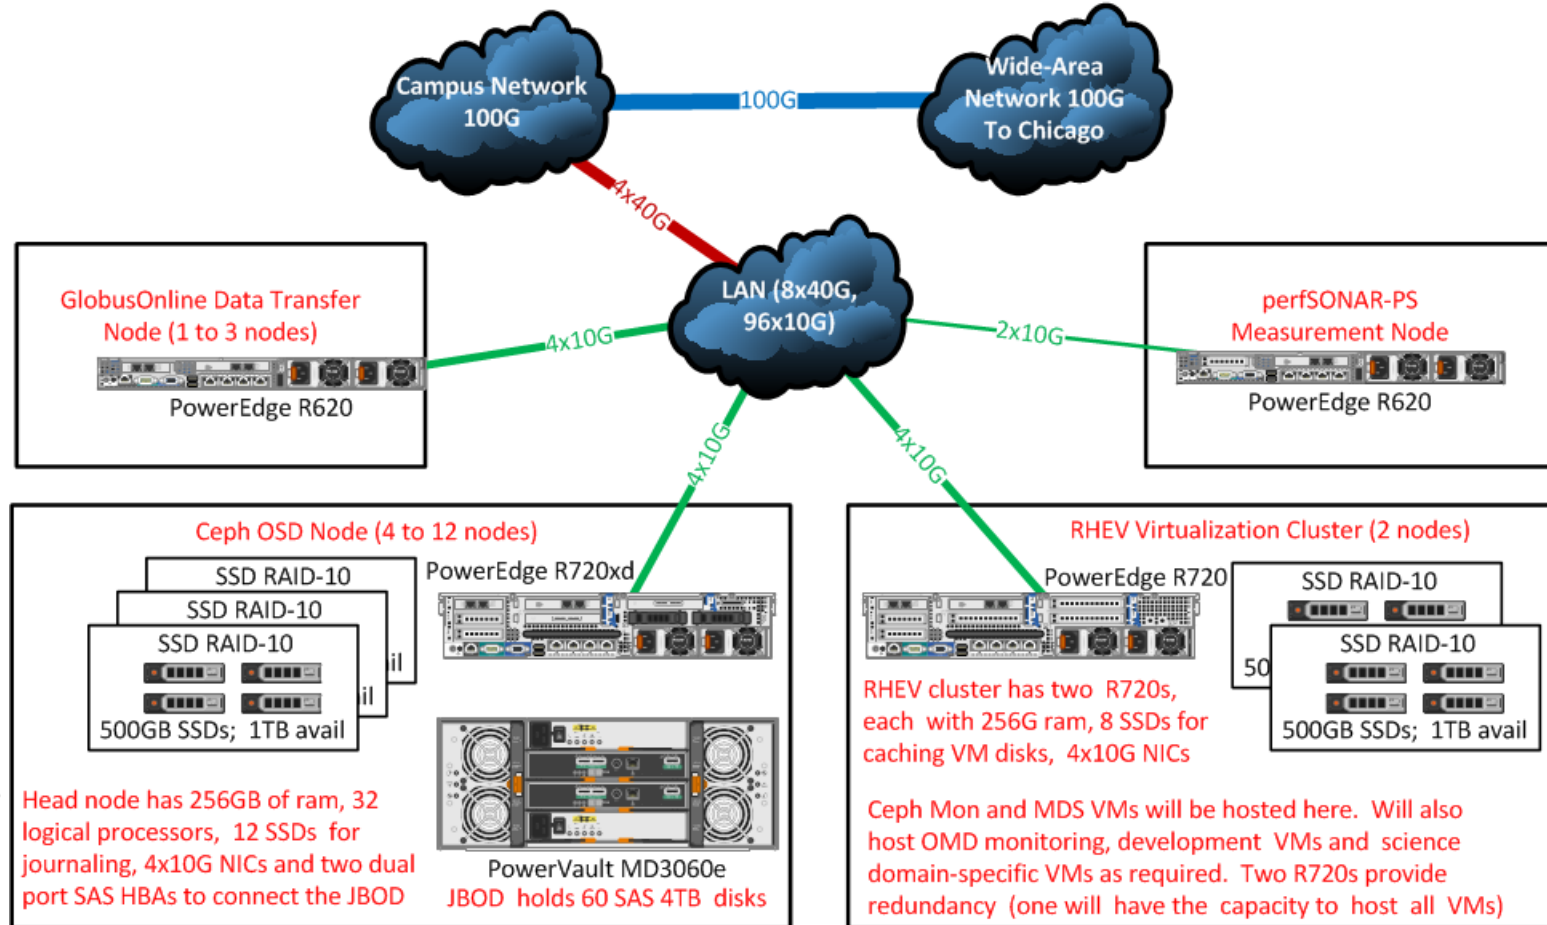
Software-Defined Storage Research

OSiRIS Data Infrastructure Building Block

NSF proposal submitted involving campus and our Tier-2

Exploring **Ceph** for future software-defined storage

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



Intends to leverage Dell “dense” storage MD3xxx (12 Gbps SAS) in JBOD mode

Summary

- New monitoring is helping us more easily find/fix issues
- CFEngine 3 in place and SL6 transition is basically complete
- Virtualization tools working well and we are close to meeting our site resiliency goals
- Network upgrade to 2x40G is nearing completion
- FUTURE: OpenStack, Condor/OS tweaks (See Brian Bockelman's talk later today)

Questions ?