



AGLT2 Site Report

Shawn McKee/University of Michigan

HEPiX Spring 2013

INFN, Bologna, Italy

April 16th, 2013

Overview of Presentation



❄ Outline

- ❑ New web pages for AGLT2
- ❑ Update on virtualization
- ❑ New Storage Hardware performance
- ❑ Provisioning system
- ❑ perfSONAR-PS Work
- ❑ Networking upgrade planning
- ❑ Conclusion

Site Summary

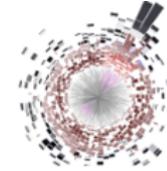


- ❄ The ATLAS Great Lakes Tier-2 (AGLT2) is a large distributed LHC Tier-2 for ATLAS that spans two locations: UM/Ann Arbor and MSU/East Lansing with roughly 50% of the storage and compute at each site
 - ❑ 4200 (4616 soon; adding 13XM420 blades) single-core job slots
 - ❑ 2 eight-core (multi-core) job slots
 - ❑ Additional 350 Tier-3 job slots (useable by Tier-2)
 - ❑ Average 8.51 HS06/job-slot
 - ❑ 3.5 Petabytes of storage
- ❄ We run most of our Tier-2 services virtualized in VMware
- ❄ We have excellent networking in place (10 GE ring to Chicaco for WAN; multi-10GE servers in LAN) with a total of ~1900 1G ports, ~350 10G ports and 8 40G ports.
- ❄ Power, space and cooling OK but at about steady-state now

New Web Pages for AGLT2



ATLAS Computing and Muon Calibration Center



Ben Meekhof put together a new website for AGLT2:

<http://www.aglt2.org>

We maintain most of the dynamic information in a set of Wiki pages

New pages group related items in tabs

The screenshot shows the AGLT2 website interface. At the top, there are navigation tabs: Home, Computing, Calibration, Projects, General, Media, People, and Wiki. Below these are sub-tabs: AGLT2 Overview, ATLAS Information, and Higgs Boson Panel. The main content area is divided into several sections:

- News:** Contains a link for "Discovery of the Higgs Boson Special Panel Discussion - Slides and Video" and a highlighted announcement: "AGLT2 demonstrates 100Gb/s network transfers at SuperComputing 2012!". A note mentions "HEPIX may be coming to UM in Fall 2013!".
- Current Statistics:** Displays "We have 7251 Condor jobs (4537 running, 2714 idle, 0 held)" and "Total Slots 4616". A link for "Job status page" is provided.
- ATLAS Computing:** Accompanied by an image of server racks, this section describes the infrastructure: "AGLT2 provides more than 4000 CPU cores and 3.5 Petabytes of storage for ATLAS physics computing. Site infrastructure services for job management, storage management, and interfacing with the ATLAS computing cloud are managed at UM and computing/storage resources are located both at UM and at MSU." It also notes that the site appears as one entity to outside users.
- ATLAS Muon System Calibration:** Includes a diagram of a tube detector system with labels for "Tube", "Preamplifier", "Discriminator", and "TDC". The text explains the calibration process: "To determine calibration compensations for ATLAS Monitored Drift Tubes a special data stream is sent to calibration centers in Michigan, Rome, and Munich." It also states that resources at AGLT2 are dedicated to MDT calibration.

The pages are much nicer for the general public and replaced a much more technical/management Wiki front page we used to direct people to.

Update on Use of Virtualization



- ❄ We are running vSphere 5.1 (since September 2012)
- ❄ We use iSCSI storage at UM and DAS at MSU
 - ❑ Working on how best to share access for iSCSI/DAS locations
 - ❑ MSU instance setup as another site
 - ❑ Using a common security and sign-on configuration between sites
- ❄ Still working on “site resiliency”. VMware base-systems and backend storage are installed at MSU
 - ❑ Exploring VMware capabilities like vSphere Replication
 - ❑ Working on overall site configuration and options (how best to both leverage relevant features & integrate with the instance at UM?)
- ❄ Goal is to have MSU capable of bringing up VMs within 1 day after the loss of the UM site
 - ❑ Services need specific planning and testing (partially complete)
 - ❑ Plan to have MSU’s VMware functional by July (SQUID soon)

New Storage Hardware Update



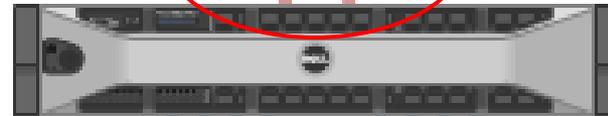
- ❄ As reported in Beijing we added at least 1.296 PB of useable space last Fall. New dense storage from Dell.
- ❄ UM added R720+2xMD3260+2xMD3060
- ❄ MSU added 2xR720 + 4xMD3260
- ❄ We were concerned with having such a large amount of storage behind a single head-node. Worry about IOPS and bandwidth / TB compared to our current R710+6xMD1200 configuration.
- ❄ We explored DDP and found it not suitable (see below)
- ❄ Using the new storage in “standard” RAID-6 mode has finally worked well. Some initial issues with kernel crashes fixed by newest “stock” kernel. Have seen 1.6 GB/sec out.

New Storage at UM



2x40G had driver issues.
Currently 2x10G

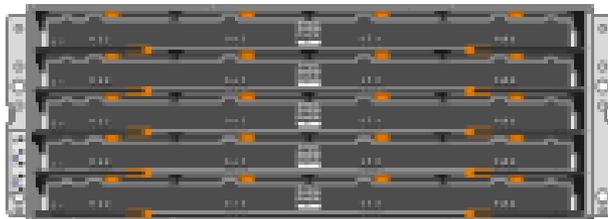
PowerEdge R720 LAG 2x40G



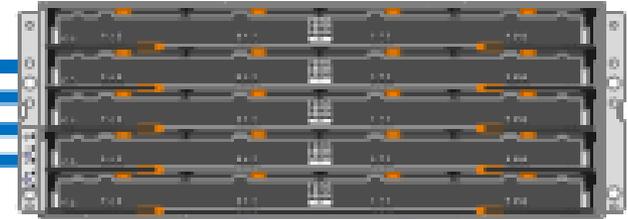
R720, dual E5-2620, 2.0 GHz, 256G 1333 MHz RAM
PCIe Gen3 (6x8,1x16) bus
4 SAS HBAs

We have 162TB/ shelf
(54TB/pool) usable for a
total of 648TB

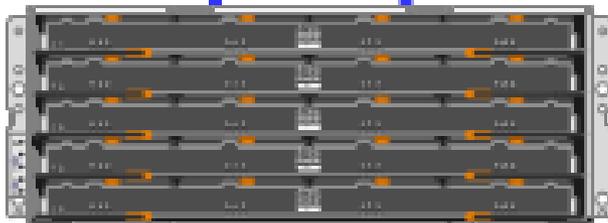
PowerVault MD3260



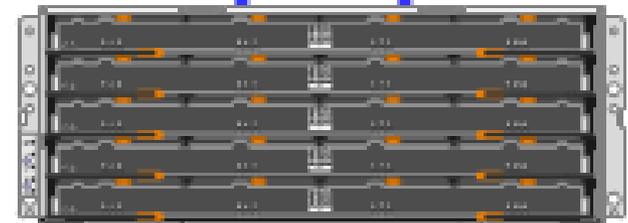
PowerVault MD3260



MD3x60 shelves hold
60x3TB NLSAS 7.2K
disks. Configuration
is 3 x 20-disk RAID6



PowerVault MD3060e



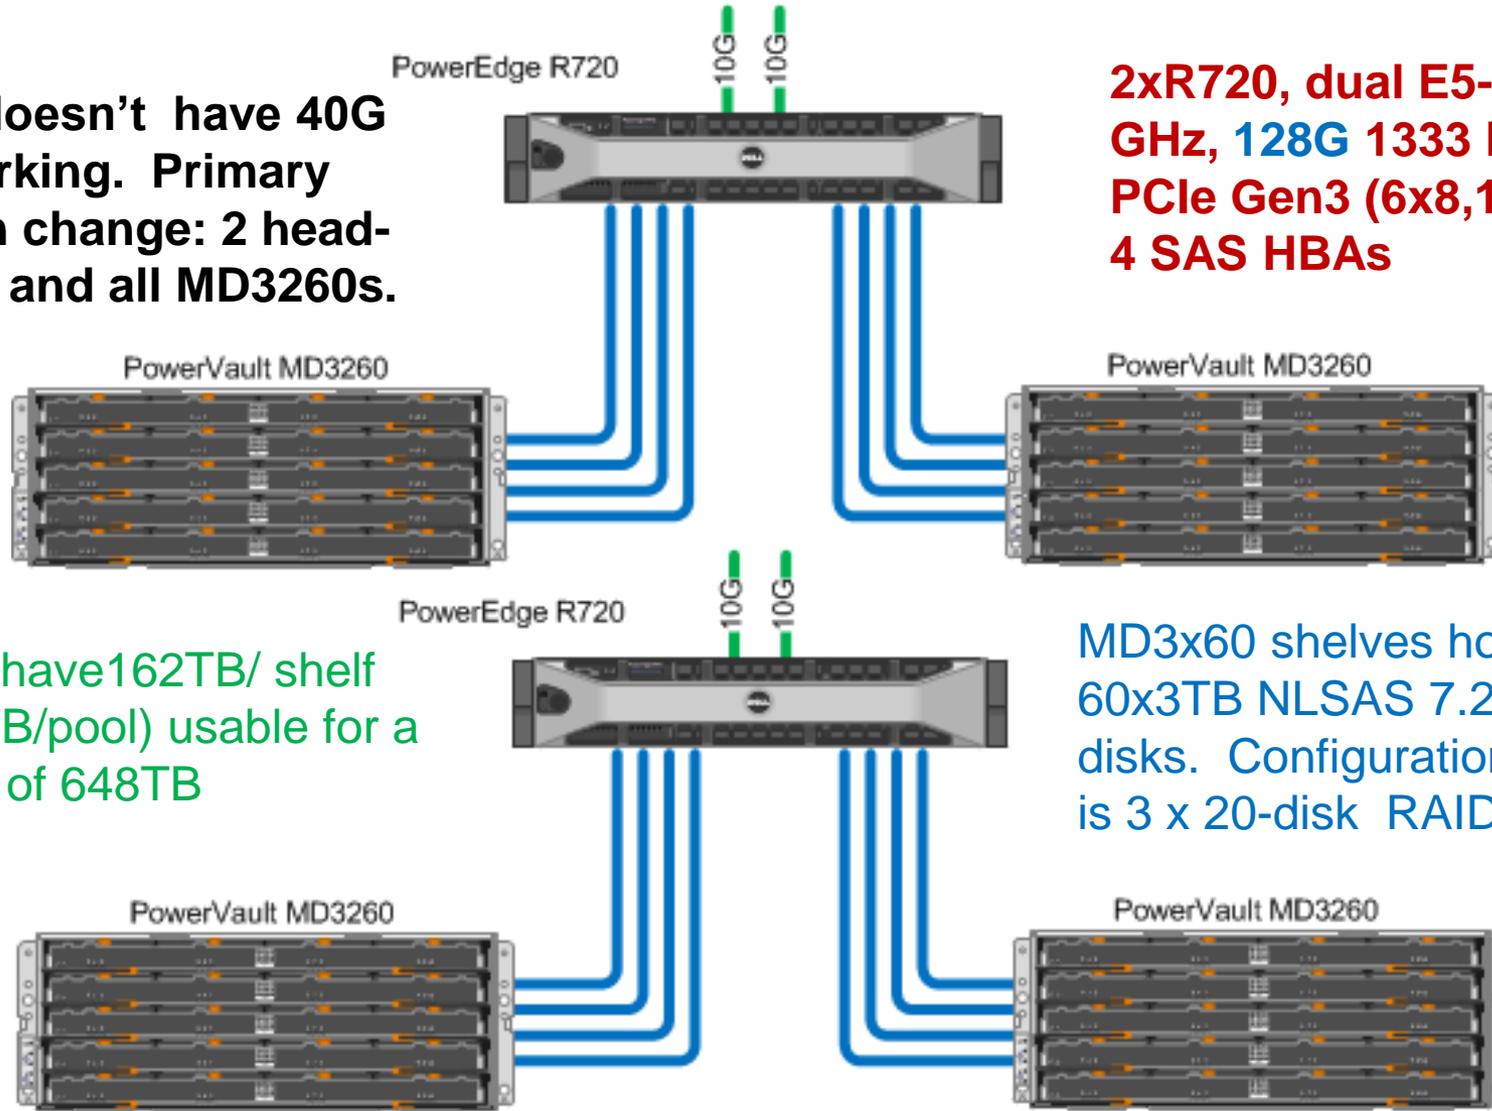
PowerVault MD3060e

New Storage at MSU



MSU doesn't have 40G networking. Primary design change: 2 head-nodes and all MD3260s.

2xR720, dual E5-2620, 2.0 GHz, 128G 1333 MHz RAM PCIe Gen3 (6x8,1x16) bus 4 SAS HBAs



We have 162TB/ shelf (54TB/pool) usable for a total of 648TB

MD3x60 shelves hold 60x3TB NLSAS 7.2K disks. Configuration is 3 x 20-disk RAID6

Testing New Storage DDP Mode



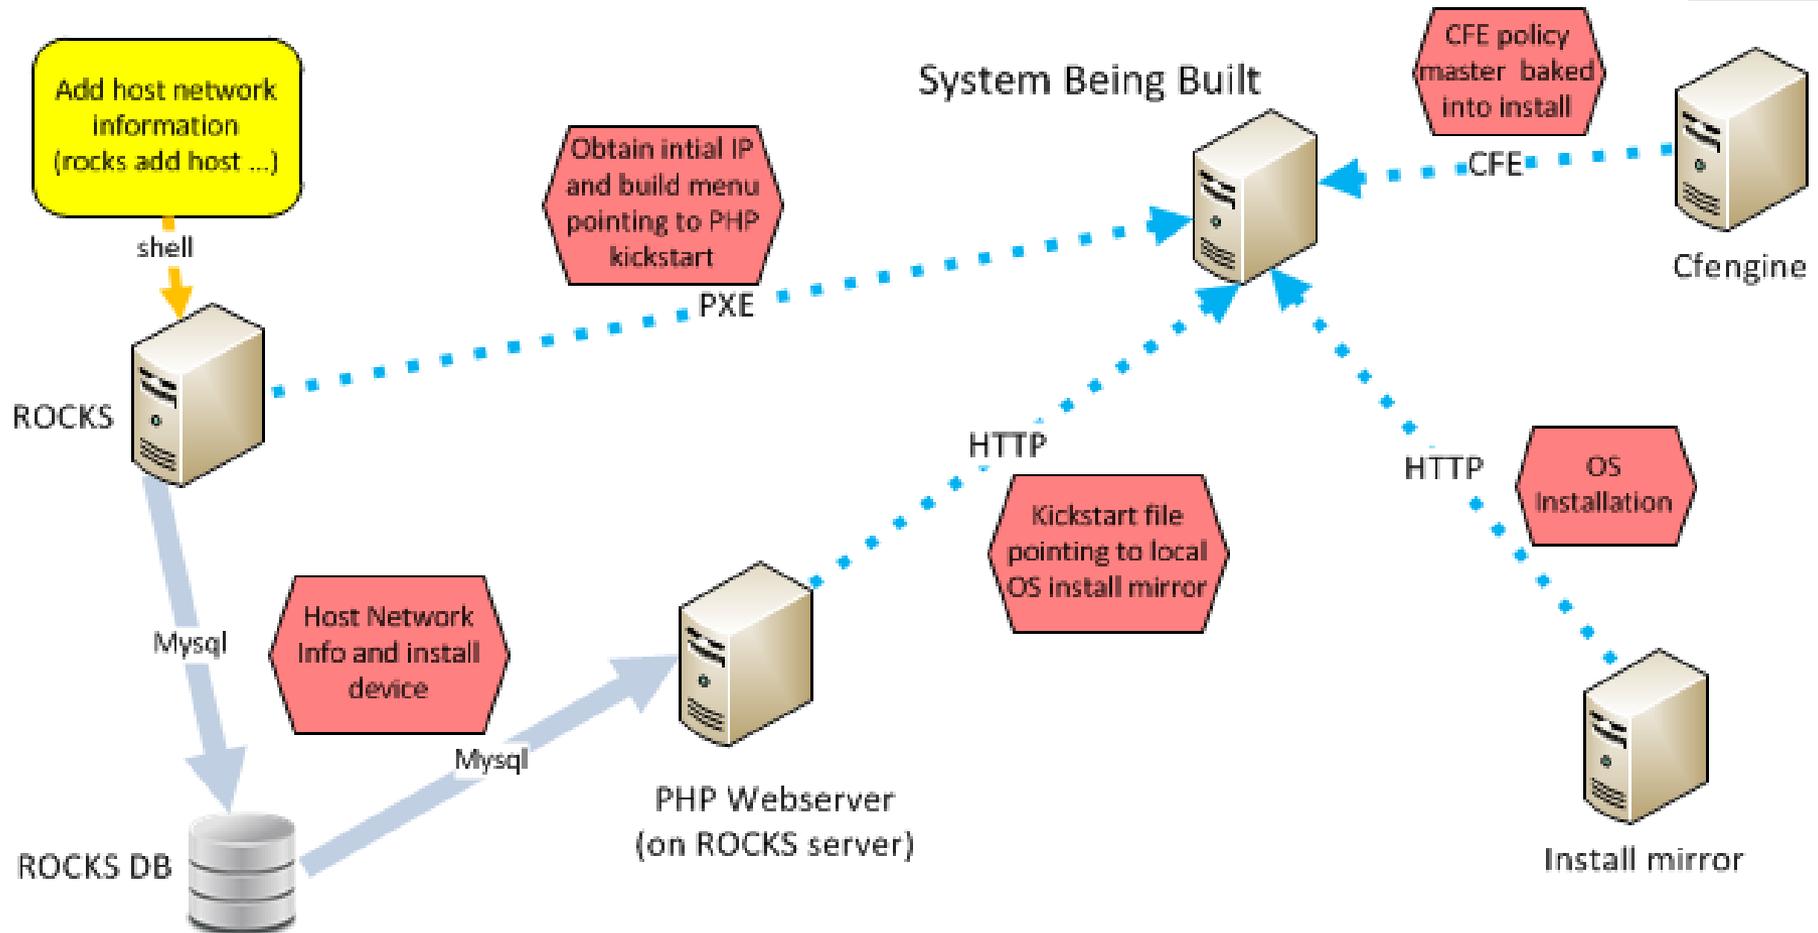
- ❄ We were very interested to test “Dynamic Disk Pools” (DDP). They provide ~RAID-6 protection and “x10 faster rebuilds”
- ❄ DDP testing results are shown on our web page at <https://www.aglt2.org/wiki/bin/view/AGLT2/Md3260Benchmark>
- ❄ Overall it didn't seem to be a good match for our use-case
 - ❑ Time to get array safe was about ½ the corresponding RAID-6 rebuild time. After array was “recovered” we still needed “rebuilding”. Total time is similar to RAID-6 total time
 - ❑ DDP ability to dynamically expand not so useful for us...we use all the space we provision.
 - ❑ DDP requires significant overhead. Extra disks are removed from storage pool losing about 7% space (4 out of 60 disks) + RAID-6 per 10 disks.

Current Provisioning System



- ❄ Service and storage node provisioning is currently handled by a combination of a **PXE menu system**, **ROCKS** build system, and a **PHP web program** that interfaces with our ROCKS database.
- ❄ We use standard ROCKS commands to create a rocks database entry which in turn sets up DHCP and PXE to point to a PXE menu allowing choices of several system types. The system is booted and kickstart is pointed at the PHP program with appropriate URL parameters for the system type. The PHP program determines network info from ROCKS DB.
- ❄ **Example: `kickstart.php?type=server&type=server-lite&os=sl6`** produces an SL6 machine with super-class "server" and more specific type of "lite server" with smaller /var and /tmp partitions.
- ❄ **Cluster compute nodes built with standard ROCKS provisioning system. We find this system too inflexible for other node types.**
- ❄ All systems configuration managed by CFEngine installed by default during build.
- ❄ **This somewhat hacked-together system allows us to utilize existing services until we transition to an overall more flexible system for provisioning (Foreman/Puppet, Cobbler).**

Diagram of AGLT2 Provisioning Setup



❄️ AGLT2 provisioning hybrid of ROCKS/PHP/CFEngine

perfSONAR-PS at AGLT2



- ❄ I will report tomorrow on WLCG and perfSONAR-PS but I wanted to update you on how we use perfSONAR-PS at AGLT2 in this site report
- ❄ The UM site has 3 sets of instances on physical machines
 - ❑ Original KOI boxes (5 years old) are now in Tier-3 area
 - ❑ Test Dell R410 using KVM to provide bandwidth and latency instance
 - ❑ New Dell R310/R610 “production” versions at Tier-2
- ❄ In addition UM has a virtualized latency instances “pinned” to each VMWare host system to verify vSwitch connectivity
- ❄ MSU also has 3 sets of instances within their Tier-2
 - ❑ Original KOI boxes at the network distribution layer
 - ❑ Production instances (Dell R310s) at access layer (near storage)
 - ❑ Additional two boxes at the AGLT2 network border
- ❄ Allows debugging of LAN as well as WAN

Planning for Detailed Stats Acquisition



- ❄ We would really like to have better visibility into our job behavior at AGLT2. Looked around a bit...
- ❄ Found TACC_STATS, a package from the Texas Advanced Computing Center, an XSEDE site, which provides detailed metrics and is being deployed XSEDE wide
https://github.com/TACCProjects/tacc_stats
- ❄ Planning to deploy this with hooks into our Condor batch system. Should provide detailed I/O, network, storage, memory and CPU stats per job and per node.

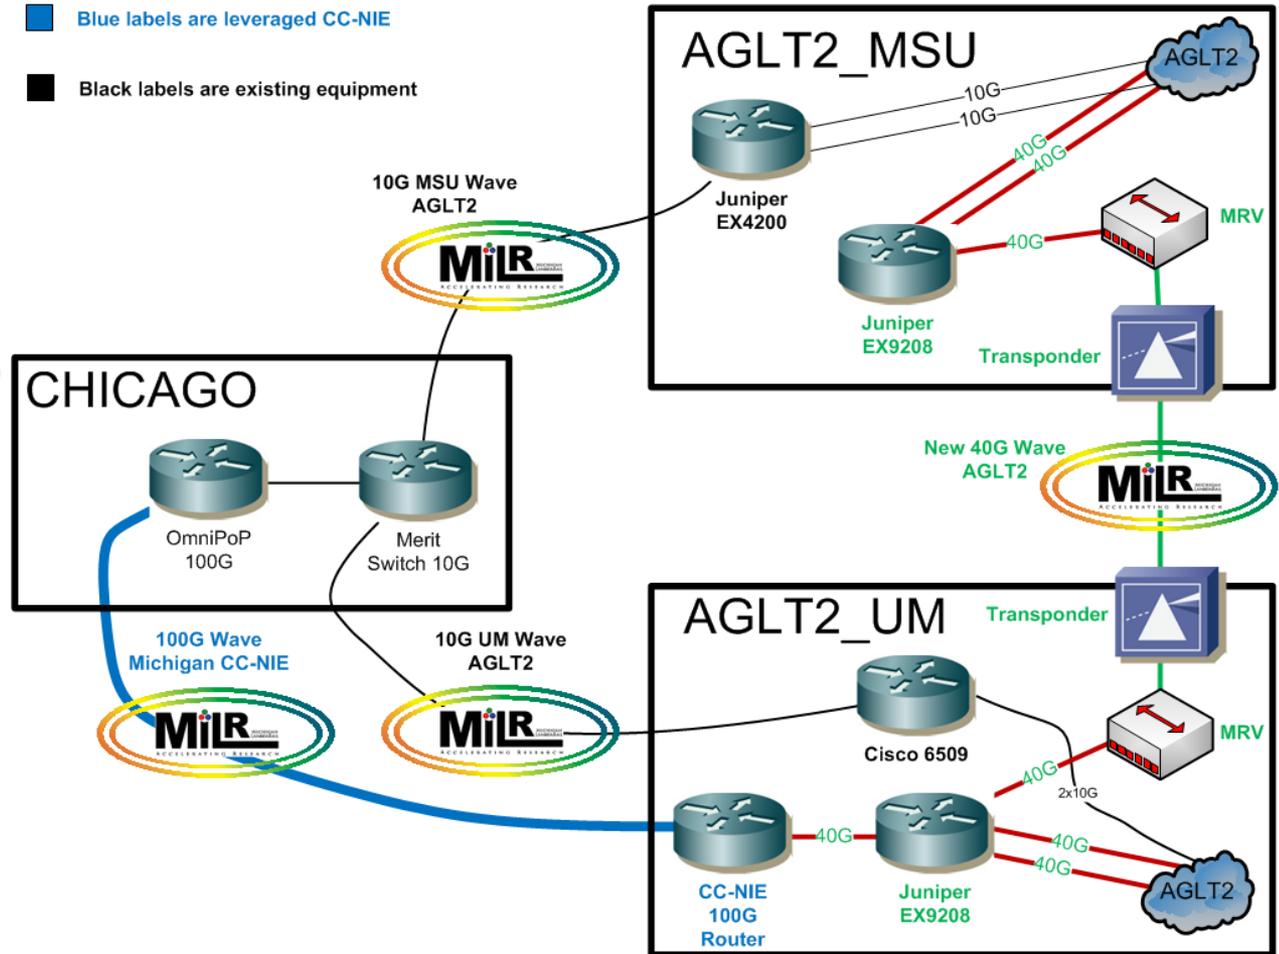
AGLT2 Planned Networking Upgrades



AGLT2 Network Upgrade Diagram

- ❄ We are exploring a possible network upgrade for AGLT2
- ❄ Current inter-site connection is frequently saturated, even after using various caching tricks in dCache
- ❄ **New 40G inter-site and shared access to 100G WAN**

- Green labels are purchased in this proposal
- Blue labels are leveraged CC-NIE
- Black labels are existing equipment



Summary



- ❄ Our Tier-2 has been running well and making progress on our “Site Resiliency” plans
- ❄ We implemented a new dense storage options from Dell that provides us with a significant amount of storage with very good performance
- ❄ Planning for next-generation networking in progress. Possible 100G access by Fall 2013. OpenFlow/SDN too.
- ❄ New effort to update our provisioning system underway...going to SL6/Foreman/Puppet this summer
- ❄ Trying to better instrument for gathering job metrics to allow detailed view into infrastructure performance.
- ❄ Looking forward to seeing everyone in Ann Arbor this Fall!



Questions / Comments?



ADDITIONAL SLIDES

Virtualization Considerations



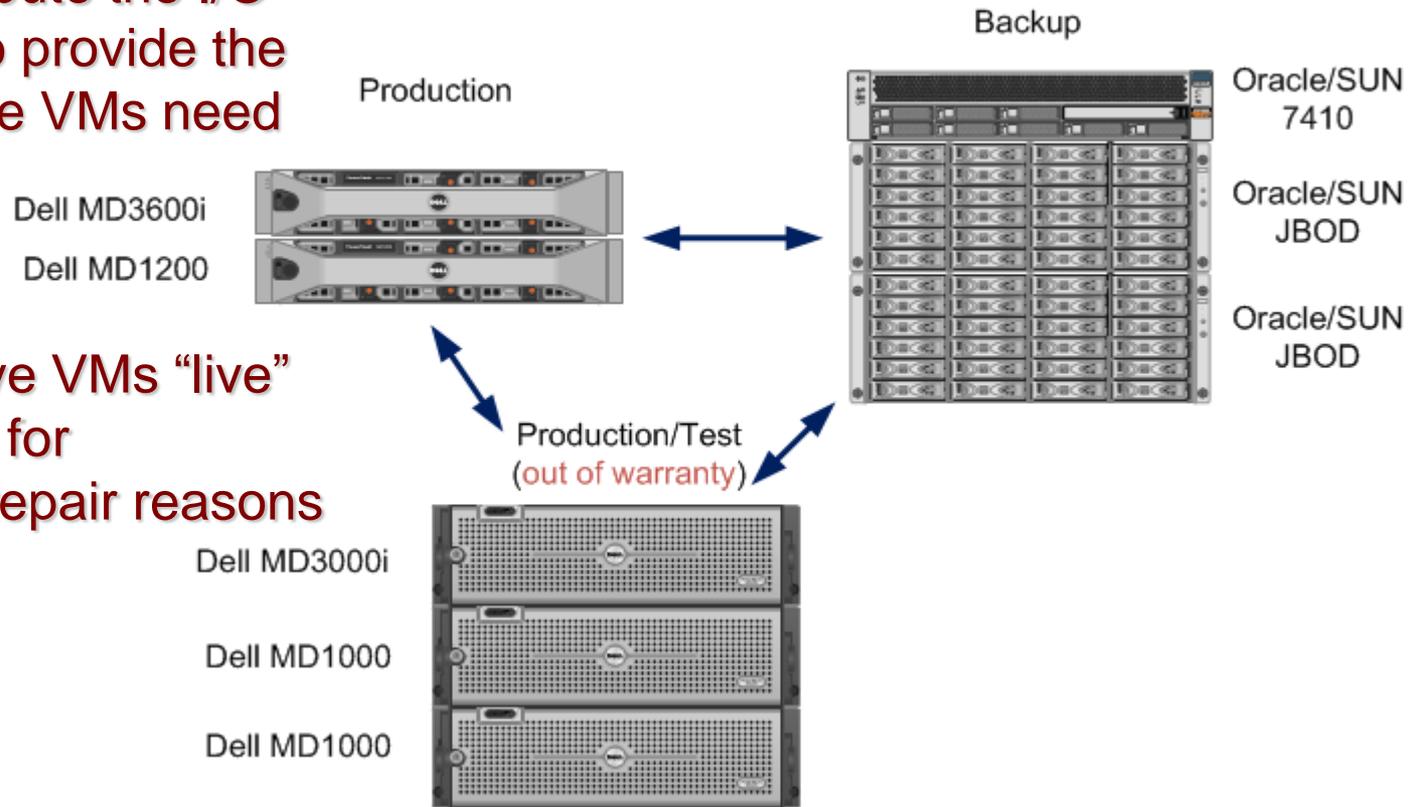
- ❄ Before deploying any VM technology you should plan out the underlying hardware layer to ensure a robust basis for whatever gets installed
- ❄ Multiple VM “servers” (to run VM images) are important for redundancy and load balancing
- ❄ Multiple, shared back-end storage is important to provide VM storage options to support advanced features
 - ❄ iSCSI or clustered filesystems recommended
- ❄ Scale hardware to planned deployment
 - ❄ Sufficient CPU/memory to allow failover (N-1 svrs)
 - ❄ Sufficient storage for range of services (IOPS+space)
 - ❄ Sufficient physical network interfaces for # of VMs
- ❄ Design for no-single point-of-failure to the extent possible

iSCSI for “backend” Live Storage Migration



❄ This set of equipment + VMware allows us to distribute the I/O load as needed to provide the level of service the VMs need

AGLT2 iSCSI Storage



❄ We can also move VMs “live” between storage for maintenance or repair reasons

Interim purchase as MSU will only allow MSU nodes to replicate VM images from UM site. This is OK for our primary goal as long as we maintain a timely replica

Storage Connectivity



❄ Increase robustness for storage by providing resiliency at

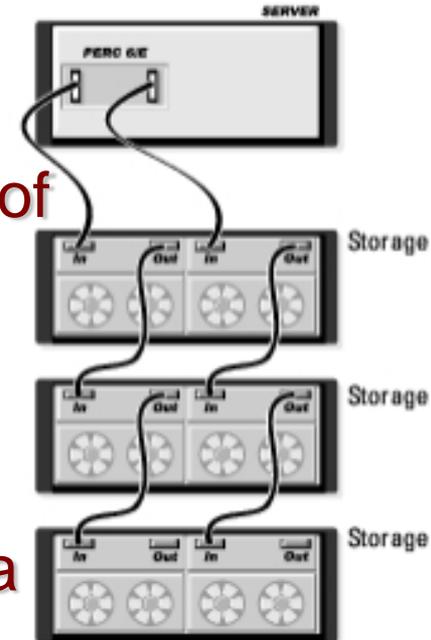
various levels:

- ❑ Network: Bonding (e.g. 802.3ad)
- ❑ Raid/SCSI redundant cabling, multipathing
- ❑ iSCSI (with redundant connections)
- ❑ Disk choices: SATA, SAS, SSD ?
- ❑ **Single-Host resiliency**: redundant power, mirrored memory, RAID OS disks, multipath controllers
- ❑ **Clustered/failover** storage servers
- ❑ Multiple copies, multiple write locations

Redundant Cabling Using Dell MD1200s



- ❄ Firmware for Dell RAID controllers allows redundant cabling of MD1200s
- ❄ MD1200 can have two EMMs, each capable of accessing all disks
- ❄ An H800 has two SAS channels
- ❄ Can now cable each channel to an EMM on a shelf. Connection shows one logical link (similar to “bond” in networking): Performance and Reliability



Storage Example: Inexpensive, Robust and Powerful



Dell has special LHC pricing available. US ATLAS Tier-2s have found the following configuration both powerful and inexpensive
Individual storage nodes have exceeded 750MB/sec on the WAN

AGLT2 V6 Storage Node
Five to Six partitions 30TB each
(Each RAID6 uses all 12 3TB disks)

