VANDERBILT UNIVERSITY

ACCRE
Advanced Computing Center for Research & Education

US CMS

Vanderbilt University Site Report
2014 OSG All Hands Meeting

Kevin L. Buterbaugh
(Kevin.Buterbaugh@vanderbilt.edu)
for the Vanderbilt Team
April 7th, 2014

A Brief Review Of Our Site

- The ACCRE cluster is a shared resource used by more than 700 researchers from 30 campus departments and four schools.
- The cluster currently has approximately 750 nodes and 7,200 processor cores.
- Each node has between 3 and 12 GB of RAM per core.
- Home directories and scratch space are stored on filesystems utilizing IBM's GPFS with a total usable capacity of 500 TB and a peak bandwidth of 13 GBytes/second.
- ACCRE Technical Staff include 10 support personnel with more than 60 years of combined experience.
- ACCRE is staffed Monday - Friday from 7:30 AM - 6 PM Central with evening / weekend support for critical issues.

Review of Grid Infrastructure

- Two Compute Elements
- Two Storage Elements
- Six GridFTP Servers
- One GUMS Server
- Two Frontier / Squid Servers
- 44 Storage Depots with 3.55 PB usable space for CMS data
- Nagios is used to monitor our status on the OSG RSV Tests and the CERN SAM tests, as well as our PhEDEx transfer rates

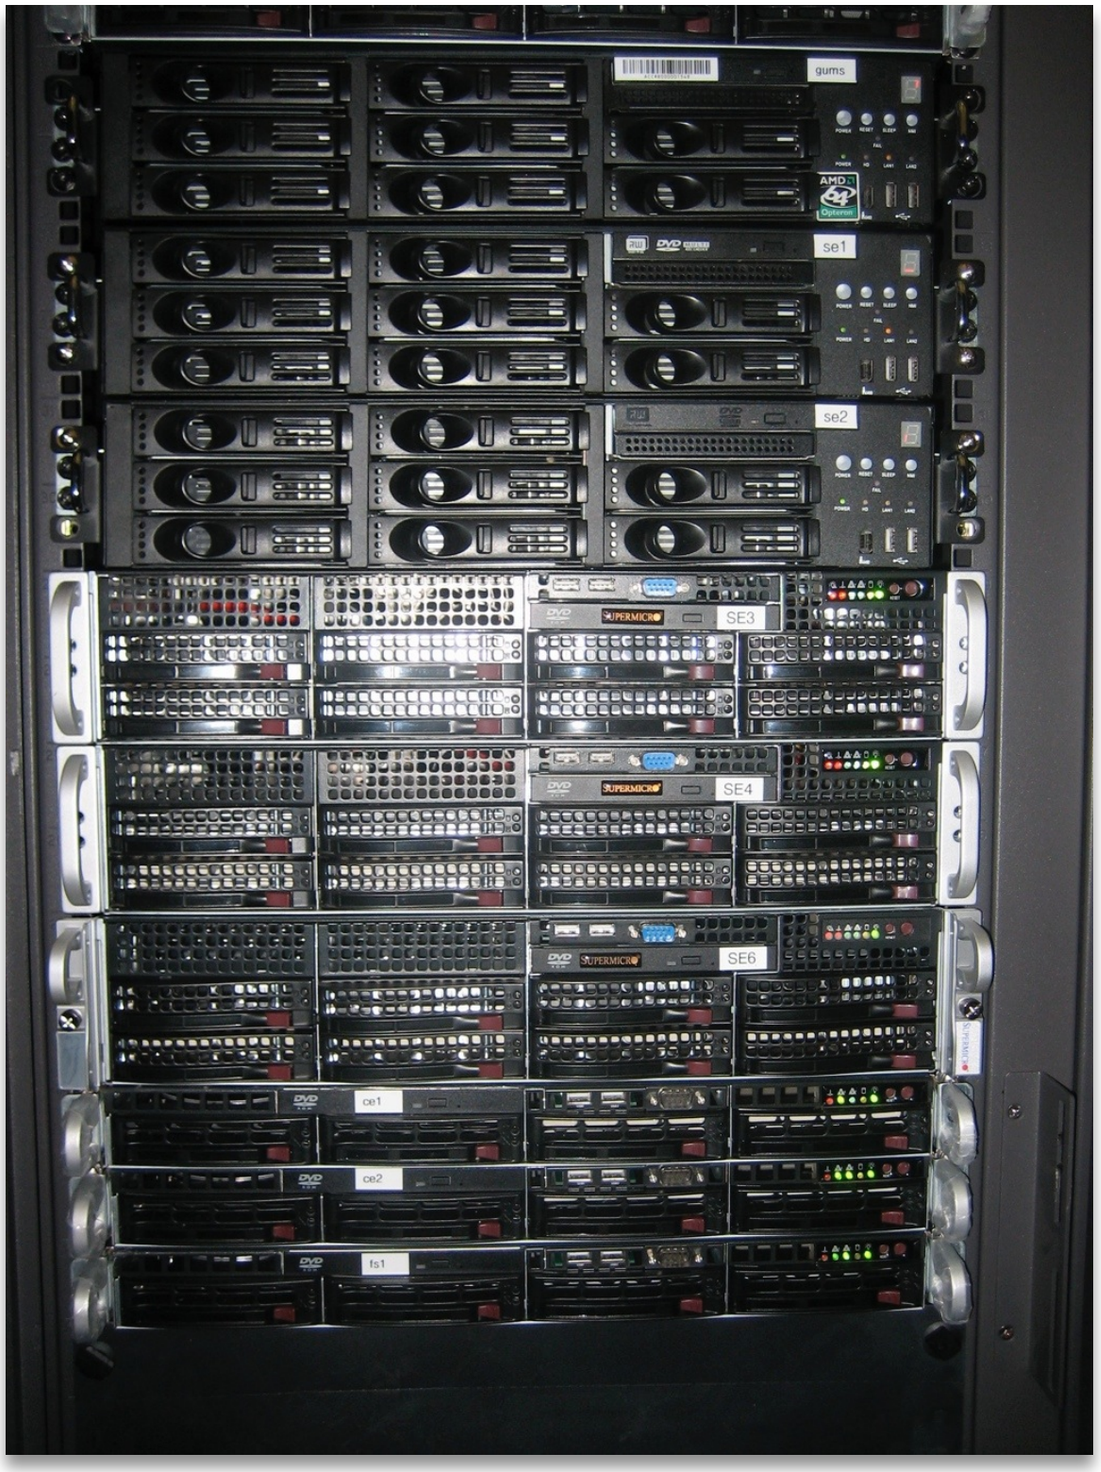

Review of Compute Nodes

- ACCRE cluster contains ~7,200 cores.
- 1368 cores purchased by / for CMS Heavy Ion (19% of the current cluster).
- CMS can currently "burst" up to 2,100 cores.

Network Architecture

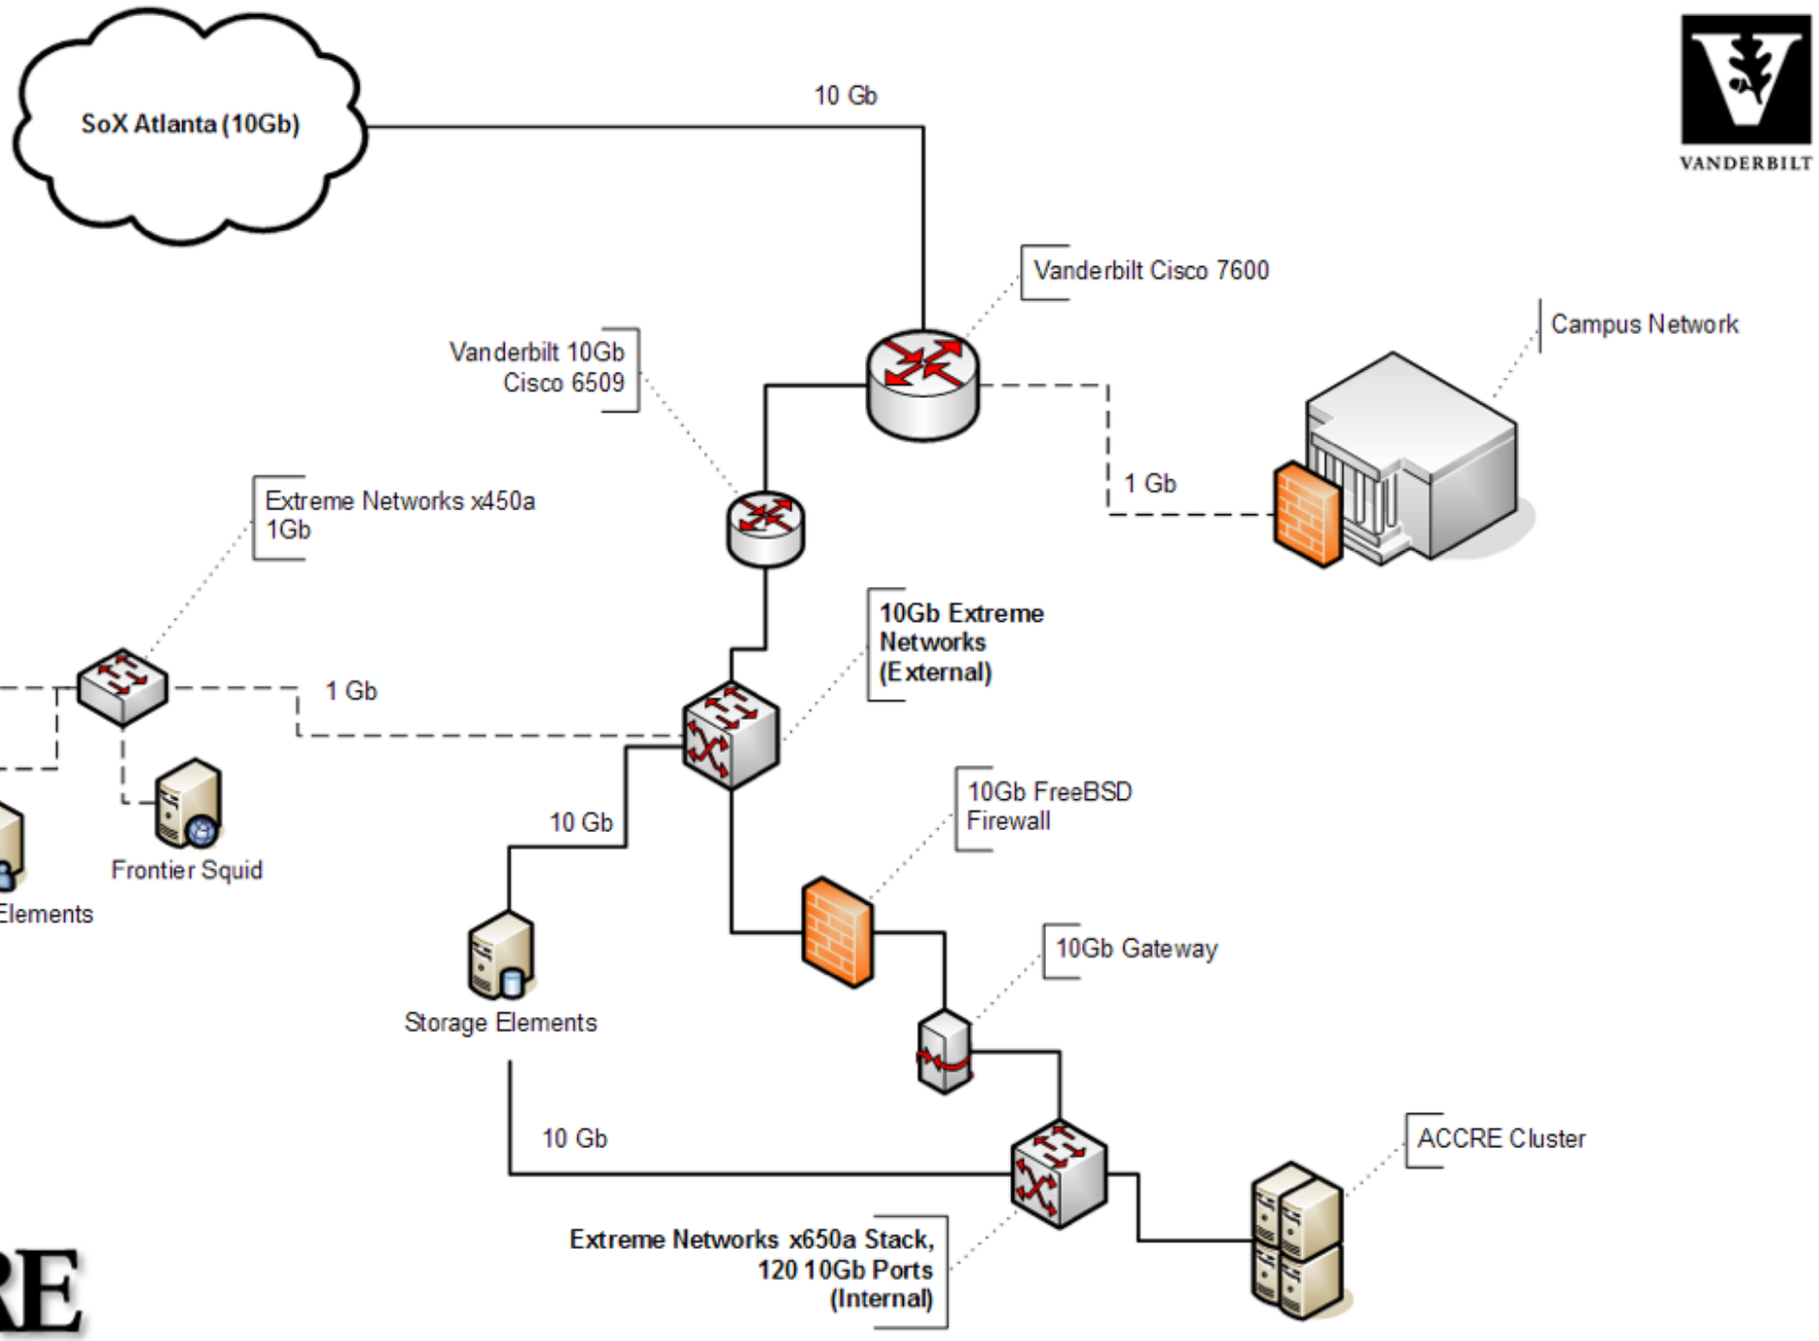

With that review out of the way…

- On to our main focus … how are we preparing for the future?
- "Prediction is difficult, especially about the future."
- A quote attributed to everyone from Neils Bohr to Yogi Berra!

What are you doing to prepare for 100 Gbps? IPV6?

- 100 Gbps
  - We're in negotiations for a 100 Gbps backhaul to SOX in Atlanta.
  - Once we have the link we will upgrade our network infrastructure as demand and funding dictates.
    - May go in phases, first to 40 Gbps, then 100 Gbps.
    - Or we may go straight to 100 Gbps if funding available.
- IPv6
  - We're ready to enable IPv6 if the need arises.

What are your capabilities for multi-core?

- All nodes in cluster are 8 core or 12 core.
- We use PBS / Moab for our cluster scheduler, which can handle multi-core and multi-node jobs.
- We would just have to configure pbs.pm to request the appropriate number of cores.
- Or is there something coming which we're not aware of which makes it more complicated than that?

What are you doing to modernize your site?

- Doing routine upgrades:
  - Replacing aging Grid infrastructure hardware.
  - Upgrading to CentOS 6 as we replace hardware.
- ACCRE is a faculty driven facility and they are guiding us on future requirements.
  - One area of interest is massively parallel processing.
    - We've had an NSF MRI award to purchase GPU nodes.
    - Working on acquiring funding for a rack of Intel Xeon Phi co-processor equipped nodes.
  - We are co-PI's on the DYNES and ANSE projects.

How is your campus computing environment changing?

- And how will your site be evolving as a result?
- The 2013 merger of ITS (University) and MIS (Medical Center) IT departments has resulted in a "one network" initiative.
- See diagram on next slide...

Vanderbilt Network Plans

Current Vanderbilt Network

Geographically Based

2 Networks

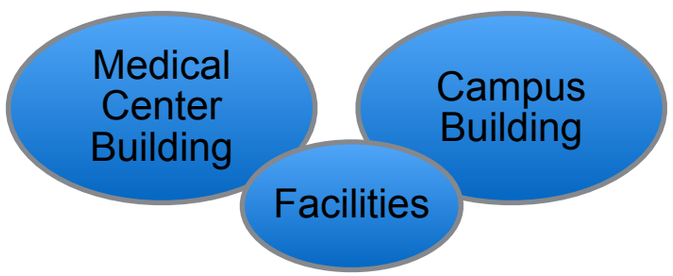

Future Vanderbilt Network Phase 1

1 Network

Wired - Manually Role Based, Broad Virtual Communities

Wireless - Geographically Based

July 2014

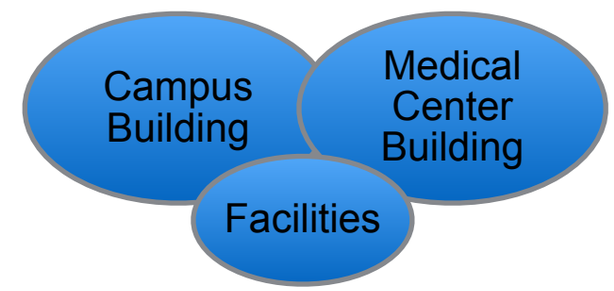

Future Vanderbilt Network Phase 2

Wired - Manually Role Based

Wireless - Auto role based by authentication

Defined Virtual Communities

January 2015

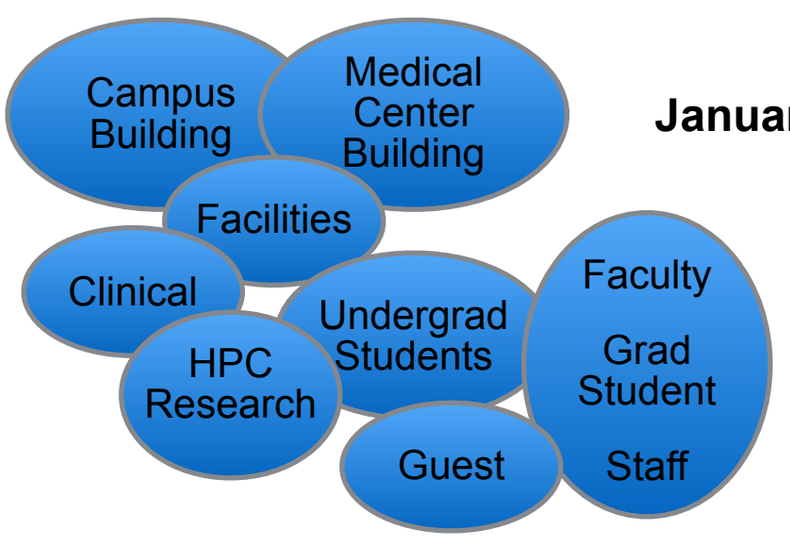

Future Vanderbilt Network Final

All Role Based - Automated by network access

July 2016

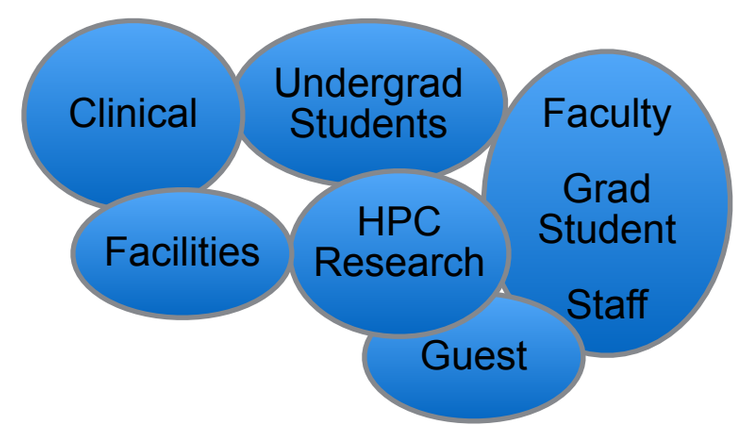

Technical improvements you are pursuing…

- …that might interest others?
- The merger of ITS and MIS into VUIT has also resulted in the research data center housing the ACCRE cluster being managed by the same people who manage the VUH data centers.
- Hospital data center managers have a completely different mindset than traditional data center managers.
- The research data center is therefore receiving numerous upgrades, beginning with electrical systems (we'll actually have a downtime in late April because of this) and continuing on to room layout, cooling, etc.