

Jefferson Lab Site Report

HEPiX Fall 2010

Cornell

Sandy Philpott

Sandy.Philpott@jlab.org

Virtual Infrastructure

VMware vSphere 4.1 with vCenter management

- 7 Hosts
 - Dual Quad Core, 96GB RAM
 - 2 hosts also contain 1TB of fast local raid storage.
 - will be used to house VMs that run mysql

150 production and 20 development Virtual Machines.

Virtual Infrastructure (2)

Storage: 4 HP Lefthand Solutions iSCSI arrays.

- Network RAID allows any 1 array to be down with no loss of availability of storage
 - 16TB total raw - using Network RAID 1 - allows 8TB useable.
- Recently licensed Cisco Nexus 1000V Switch software.
 - Replaces vmware's virtual switch
 - Runs Cisco NX-OS - allows same management and monitoring techniques as physical switches.
 - Will soon begin project to virtualize all of MIS/BSN application servers.
 - Some will be new server builds, but most will be converted (P2V).

Power Management

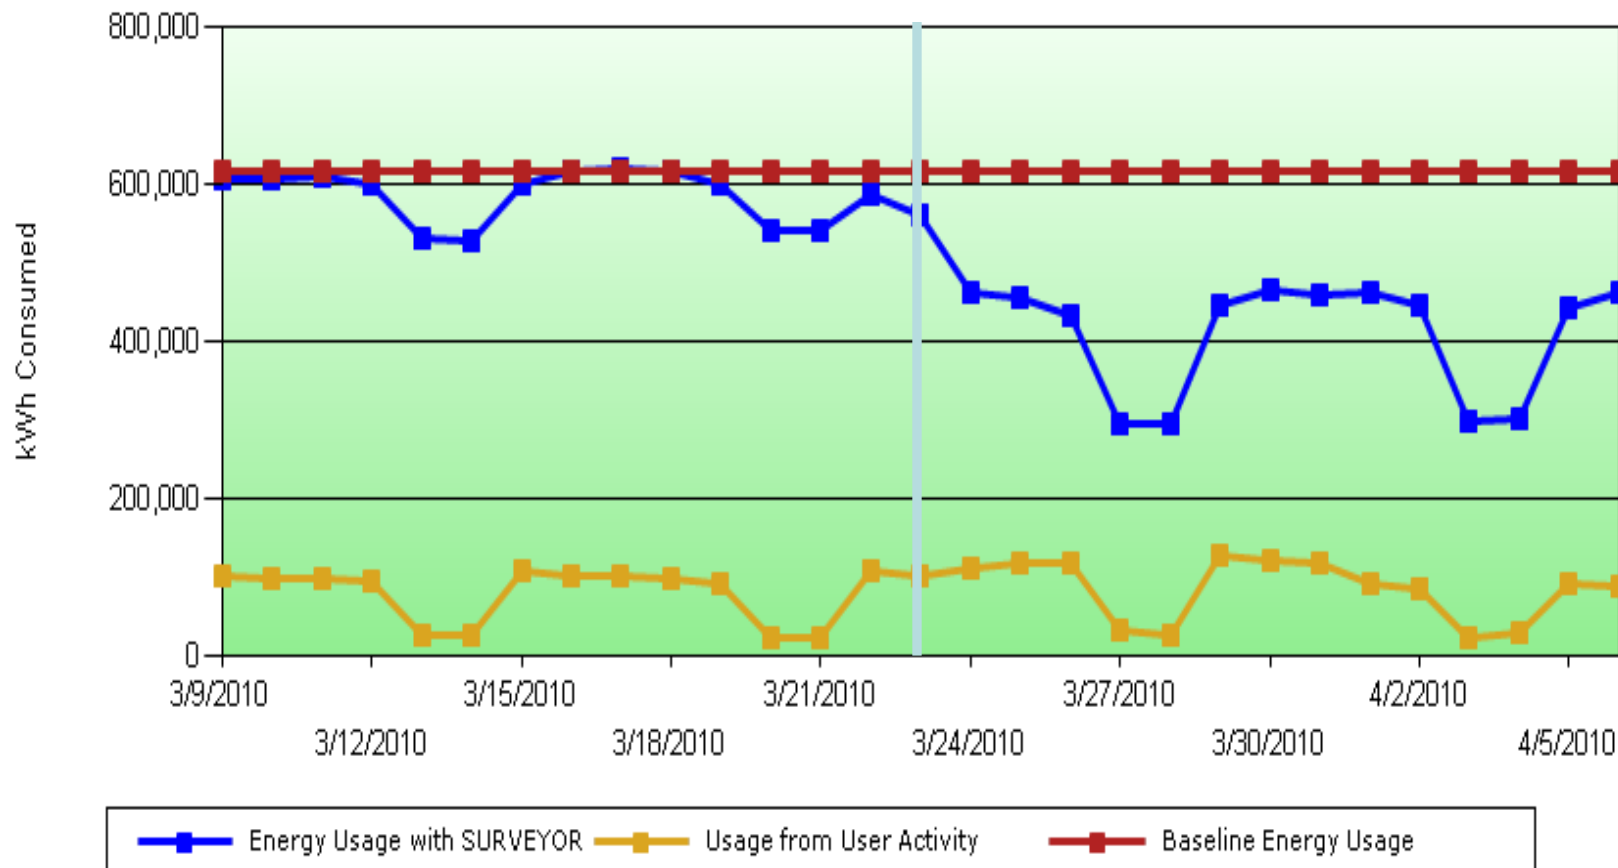
Surveyor from Verdiem to manage power consumption for Windows desktops computers

- can schedule systems to go to standby mode
- easy for scheduling global wake up weekday mornings, for patch install or for virus scan
- Can create a group and assign different policy for different groups (i.e. like Exempted Group)
- Has a builtin Wake-on-lan feature, but JLAB is using home grown solution

Power Management was turned on 3/23/2010 sitewide

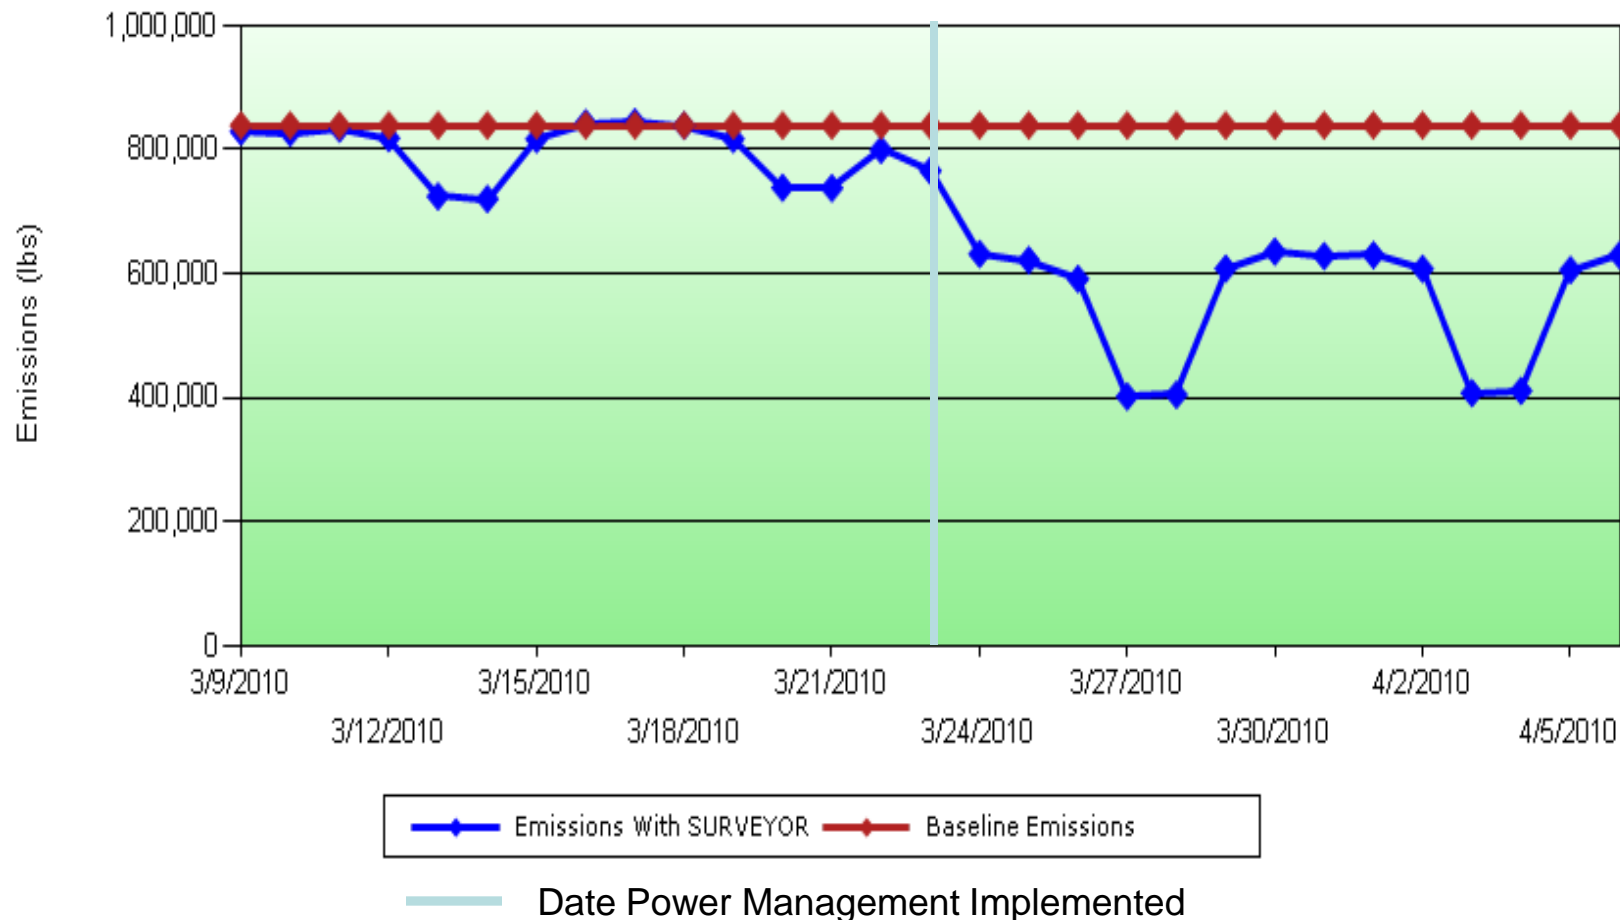
- Attached graph shows power consumption of two weeks before and two weeks after deployment.

Annualized Daily Power Usage



— Date Power Management Implemented

Annualized daily Greenhouse Gas Emissions



Identity Management / Single Sign on

Participating in the Science Identity Federation, ESnet

Planning on joining InCommon, a large educational identity federation

- authenticate web services using credentials from other federation member institutions
- Shibboleth pilot SSO server up and running
 - Password and Certificate-based authentication for JLab credentials
 - Plan to expand the certificate-based authentication (a Shibboleth add-on produced by SWITCH) to perform smartcard based authentication
 - Many applications have modules available to allow Shibboleth/SAML integration -- including Drupal and MediaWiki

Web Services

- – Drupal
 - now in use for Central Computing website
 - evaluating for primary www.jlab.org website
 - identify good core configuration to meet the needs of multiple groups
 - Seeking methods for implementing common applications like logbooks, notebooks, wikis, etc
 - Plan to develop a standard deployment

High Performance Computing

Theoretical Physics, Lattice QCD

ARRA clusters: \$4.9M, 380KW

- 10g: 50 SuperMicros from KOI Computing
 - dual quad core 2.53 GHz Westmere CPUs
 - 48 GB memory
 - 32 quad C2050, QDR Infiniband
 - 18 quad GTX-480 GPUs, SDR Infiniband
- 10q: 224 SuperMicros from KOI Computing
 - dual quad core 2.53 GHz Westmere CPUs
 - 24 GB memory
 - QDR Infiniband

High Performance Computing (2)

ARRA Clusters (cont)

- 9q cluster: 320 Dell PowerEdge R410s
 - Dual quad core 2.4GHz Nehalems
 - 24GB RAM
 - QDR Infiniband (40 Gb/s)
 - 6 sets of 32 nodes, 1 set of 128 nodes
 - 54 nodes with single GTX-285
 - 10 nodes with single C2050
- 9g: 64 SuperMicros from KOI Computing
 - Dual quad core 2.4GHz Nehalems
 - 48GB RAM
 - quad GPU nodes
 - Half GTX-285s, half GTX-480s

High Performance Computing (3)

- Storage: Lustre

- 14 SuperMicro OSSs from AMAX

- 3ware 9650 24 port controller

- 24 * 1TB disks

- MDS from Penguin Computing

- RAID10

- 6 * 1TB disks

Upgrading for redundancy; hardware is in place

Planning another storage purchase soon...

- 32 * 2TB disks

- controller?? LSI MegaRAID?

Details in Thursday's talk...

Offline Data Analysis Cluster

- Auger software from JLab, using Maui/PBS
- 75 batch farm nodes, mostly dual quad core Nehalems with 24GB RAM; CentOS 5.3 64-bit
- 25 nodes still running Fedora 8 32-bit
 - Decommissioning by year-end

JASMine, Mass Storage

IBM TS3500 tape library

8 LTO-4 drives, 4400 slots

Fully operational, >4PB; 10% capacity left

- Eject, or expand

4 LTO-5 drives; Infiniband data movers

beta testers in the spring

purchased, put into production

ongoing problems, IBM is working to resolve

drives “ate” tapes; spooled off reel

seek errors

4 brand new replacement drives being installed today

Plan to move cache files (data already in silo, now on disk) into Lustre

Upcoming

- 32-bit data analysis farm nodes to decommission
- Lustre IB/NFS gateway for farm access
- LTO-5 problem resolution
- Lustre storage upgrades
- Accelerator upgrade 6 -> 12 GeV by 2015