

Computing Division

Scientific Computing Services

Unix Town Hall Meeting

Yemi Adesanya, May 12, 2016



Unix Town Hall Meeting



Objectives:

- Communication
- Collaboration

Join our mailing list: unix-community@slac.stanford.edu

email to: listserv@slac.stanford.edu

subscribe unix-community

New Scientific Computing Services web portal

SLAC

<https://internal.slac.stanford.edu/computing/scientific-computing-services>

- Discover our Services, Resources, Expertise
- Up-to-date web documentation targeted at SLAC Unix users
- From a User's perspective: What do I need from Computing Division to get the science done?
- Included in the orientation process/docs for new employees?

Scientific Computing Services



Scientific Computing Services (confluence) page

<https://confluence.slac.stanford.edu/display/SCSPub/Scientific+Computing+Services+Home>

New web page under development

<https://internal.slac.stanford.edu/computing/scientific-computing-services>

unix-admin@slac.stanford.edu
support/questions

yemi@slac.stanford.edu
650-926-2863

Agenda:

- Computing Partnerships
- Unix Platform
- Storage & Data Management
- HPC in the Cloud
- Cyber Security update
- Questions/Discussion

Computing Partnerships

Scientific Computing Services

Yemi Adesanya, May 12, 2016



- New SLAC CS department announced on 4/1/2016
- Professor Alex Aiken will lead the group
- Support the scientists who are pushing the limits of data analysis and simulation
- The SLAC computing challenges will drive the CS agenda
- Perfect opportunity for SCS to provide CS researchers with services for development and testing
- Alex has already expressed interest in our shared compute farms

Zoox Compute Cluster and Storage

- SCS will deploy and manage Zoox Cluster and storage
- Plans for 10 water-cooled racks in building 50 datacenter
- Networking upgrades for building 82 Firehouse and bandwidth to offsite headquarters
- Evaluating GPUs for Machine Learning applications
- Flash-based storage solutions under consideration
- John-David “JD” Castillo – Zoox IT Director
- Amedeo Perazzo (LCLS, LCLSII) will participate in design and testing
- Neal Adams – Project Manager

Unix Platform

Scientific Computing Services

Karl Amrhein, May 12, 2016



Red Hat Enterprise Linux 7 Roll Out Plan

- Centrally Managed via Chef Configuration Management



1. Servers

- OpenStack (controllers & hypervisors)
- Zoox GPU (RHEL7 running Ubuntu docker containers)



2. Virtual Machines (running on OpenStack or VMware)

- Centrally Managed RHEL 7 VMs available via self-service

3. Interactive Login

- `rhel7.slac.stanford.edu` (available for test)

4. LSF Batch Nodes

- test queues available (available for test)

5. Desktops

- soon available for test

next

RHEL 7 on Desktop or OpenStack Virtual Machine ?

- Get your own RHEL 7 VM in OpenStack Cluster
 - self-service via Horizon Web Portal, or REST API
 - or email unix-admin and we'll create it
- Do your testing and development inside your own centrally managed virtual machine which is identical to production environment (servers, batch nodes, etc.)
- Use your Desktop for personal productivity, not development and testing
- VMs can be destroyed and recreated easily, a desktop configuration cannot be an exact match for production

Scope of Red Hat Enterprise Linux 7 Desktop Support

SLAC

- Desktop team provides Level 1 support for a clearly defined set of Red Hat applications
- Examples include email, browser, ssh terminals, Office Suite, etc.
- Final list of centrally supported RHEL 7 Desktop applications is under development
- Where unique and custom environments are required (eg, 3rd party non-Red Hat device driver to manage a tile display), users/groups have management and support responsibilities

- A modern, supported solution for displaying remote Unix applications on your desktop or laptop.
- Can connect from an offsite non-SLAC network
- Can use any standard web browser (Firefox, Chrome, IE, Safari)
- Or you can download and use a client application for Mac, Windows, or Linux
- Configured in a cluster to provide redundancy, performance, and can be scaled out to meet capacity demands.

Questions?

Storage & Data Management Scientific Computing Services

Lance Nakata, May 12, 2016




Storage-as-a-Service (StaaS)

- Shared NFS filesystem for moderate performance needs
- \$100/TB/year pricing for groups needing 10's of TBs
- Service in production; charges beginning in FY17
- Most NetApps already moved to StaaS due to support phase-out. Groups will need to budget for this. We will provide estimates in June to those affected.

Storage-as-a-Service (StaaS, cont'd)

- 250TB allocated out of 320TB (78%); 81TB in use (25%)
- AD: 14%, Comm: 6%, OCIO: 27% (some temp space), SCI: 23%, Other: 8%, Available: 22% (as of 5/10/2016)

Allocation



Tape Storage

- Upgrade from 1TB to 8TB tape drives
 - 8TB drives increase tape library capacity to 100PB
 - 5TB and 8TB tape drives use the same media
 - Would decrease tape purchase cost by 37.5% vs. 5TB tapes
 - Still need to find funding
- Retire unfunded astore/mstore service
 - HPSSfs and GPFS HSM are possible solutions that provide NFS-like interface
 - May require some form of charge-back unless SLAC-funded

Storage Notes and Futures

- Continue work on automated GPFS disk-to-tape file migration. Appears to require redundant hardware configuration.
- Evaluate new storage building block hardware for performance and cost savings
- Spectrum Scale/GPFS appliances appear to be too expensive at this time
- Look at local object storage as a possible disk tier
- Cloud storage is still expensive (disk & bandwidth costs)

Questions?

HPC in the Cloud

Scientific Computing Services

Yemi Adesanya, May 12, 2016



A quote from a Cray exec. at SC15 conference:

“The future of HPC is partly cloudy”

Cloud will not replace our entire datacenter (today)

- Why Cloud?
 - Scalability and flexibility
 - Use it when you need it for bursty workloads
- Contributing factors for SLAC cloud services
 - Underutilized static, bare-metal compute clusters
 - Shared clusters averaging 87% ✓
 - Private clusters may be as low as 5% ✗
 - Funding agencies may question requests for hardware purchases
 - “Can the researcher use DOE compute center, NERSC, AWS, etc.”

Flexible Onsite Clusters

- Common compute hardware shared across the Science areas
- Shared hardware = higher overall utilization
- Provisioning virtualized hosts to meet demand
- OpenStack for high-throughput batch compute
 - New “deft” nodes coming to LSF shared queues
- OpenStack self-service for private cloud VMs
 - RDO replaces Nebula
- VMware for critical science infrastructure (high availability)
- Low-latency, fast local storage for data-intensive workloads
- Hardware leasing options for lifecycle management
- Bare-metal compute clusters only when necessary
(Specialized hardware or Interconnect)

Interfacing Public Cloud

- Traditional SLAC HPC was not designed around Cloud Infrastructure
- We must provide seamless interfaces to encourage adoption of cloud
- Networking:
 - VPN between EC2 and SLAC – VMs with SLAC IPs
 - AWS node with SLAC Unix accounts, NFS, AFS (fully tailored machine)
- Batch Job Scheduler:
 - Cloud VMs configured as LSF batch nodes
 - On-demand bursting: launch additional VMs in cloud when # pending jobs exceeds a threshold
 - User could declare a job 'cloud compatible' at submission time

Getting ready for Cloud

Can your compute run in the cloud?

- Is your compute bound to local storage (AFS, NFS, etc)?
- How much data is ingested and output?
- Can you tolerate reduced I/O performance?
- Do you need special interconnect or co-processor?

Leverage container technology for portability

- Bundle runtime environment into a container image
- Use the same container image at different compute centers
- Run on a minimally configured Linux cloud instance
- Beta Docker container feature in LSF

Thoughts on funding

- There must be some birthright/trial service
 - All SLAC users get some minimum entitlement
- Goal is cost recovery of hardware (leasing) and licensing
 - We are a non-profit!
- Avoid any chargeback model with excessive accounting overhead
 - Chargeback per-group or per-project, not per individual
 - Offer a select number of service tiers or bundles
 - Quota based, not actual utilization minutes
- Any new service funding model will have to be reviewed with senior management and business managers

Questions?