Computing Division
Unix Town Hall

Yemi Adesanya, Associate Director Scientific Computing Services

June 28th, 2018

Unix Town Hall Meeting

SLAC

Objectives:

- Communication
- Collaboration

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

email to: listserv@slac.stanford.edu

subscribe unix-community

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

Unix Town Hall Meeting

SLAC

Agenda:

- Conferences & Training
- Storage & Data Management
- CentOS7/RHEL7 Update
- Linux Desktop Support
- CPU Hardware Refresh
- Singularity Containers for Batch
- GPU compute & Jupyter
- Future of AFS/YFS
- Virtualization as a Service
- Cyber Security Updates
- Questions/Discussion

Conferences and Training

- CHEP 2018
  - July 9th-13th Sofia, Bulgaria
  - http://chep2018.org
- HEPiX
  - October 8th-12th Barcelona, Spain
  - https://www.hepix.org
- SC18
  - November 11th-16th Dallas, TX, USA
  - Joint Stanford/SLAC booth
  - https://sc18.supercomputing.org
- LISA (Large Installation System Administration)
  - October 29th-31st Nashville, TN, USA
  - https://www.usenix.org/conference/lisa18

Storage & Data Management

Lance Nakata

June 28, 2018, Unix Town Hall Meeting

Storage Update

SLAC

- Storage as a Service (StaaS)
  - Price is $65/TB/yr for FY18
  - StaaS data capacity increased from 320TB to 480TB
  - Limited SSD data area created for small block, random access I/O; contact unix-admin for more info
  - GPFS version upgrade from 3.5 to 4.1.1
  - Plan: Upgrade from 4.1.1 to 4.2.3 and then consider enabling Local Read-Only Cache (LROC) for enhanced read performance
  - Plan: Investigate possible use of a Spectrum Scale appliance for overall better performance and manageability

- GPFS 3.5 to Spectrum Scale 4.1.1 upgrades ongoing
  - Only two clusters still requiring upgrading
  - Two-step upgrade process from 3.5 to 4.1.1 then later to 4.2.3 to reduce or eliminate any downtimes

Storage Update (2)

SLAC

- Spectrum Scale Licensing
  - Evaluating the practicality of capacity vs. socket licenses
  - Advantage: capacity-based licensing includes client licenses

- HPSS upgraded from 7.4.3 to 7.5.1
  - 7.4.3 is going end of life 9/30/2018
  - New Feature: Tape Ordered Recall could reduce tape mounts when staging files, thereby decreasing wait times in some cases

- tsm1 tape backup server upgrade (Q4FY18)
  - SSDs for increased database performance
  - Move from 1TB to 5TB tape drives
  - Likely target for Spectrum Scale HSM tape usage
  - Will evaluate deduplication and offsite storage capabilities

Storage and Solaris End-Of-Life

SLAC

- End-Of-Life = No longer supported by vendor. EOL hardware:
  - Solaris SPARC storage (e.g., “sulkys”)
  - Sun Thumpers/Thors (e.g., “kans, wains”)
  - LSI Engenio disk arrays (Fermi, KIPAC, SIMES)

- Regular Solaris 10 support ended 1/31/2018. Hardware phaseout will continue through 2018.

- Working on contingency plans for unsupported hardware that fails. Many old arrays and storage servers will move to StaaS if the data must be retained on disk.

- Spectrum Scale/GPFS running on RHEL is the current supported storage platform.

SLAC

Questions?

CentOS7/RHEL7 Update

Karl Amrhein

June 28, 2018, Unix Town Hall Meeting

RHEL lifecycle support

SLAC

Red Hat Enterprise Linux (RHEL)

https://access.redhat.com/support/policy/updates/errata

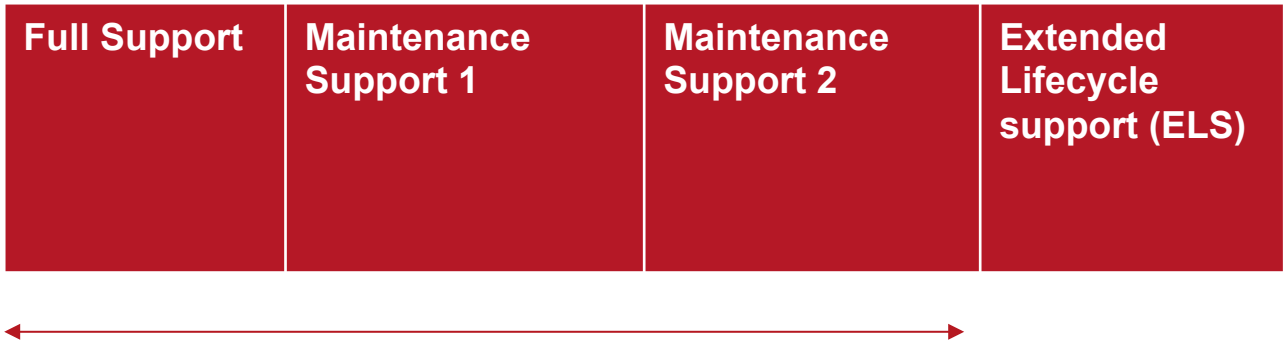

10 year lifecycle for each major version

RHEL lifecycle support

SLAC

- RHEL 5.11
  - 2007 - 2017
  - Currently in: Extended Lifecycle Support (ELS)
  - Any remaining RHEL 5.11 hosts need to be retired or updated
- RHEL 6.10
  - 2010 - 2020
  - Currently in: Maintenance Support 2
- RHEL 7.5
  - 2014 – 2024
  - Currently in: Full Support

Red Hat and CentOS relationship

SLAC

- Red Hat and CentOS joined forces in January 2014
- CentOS remains independent from Red Hat
- The joining of forces strengthens the CentOS community and facilitates the CentOS build process since Red Hat is directly involved
- SLAC has benefited from Red Hat vendor support since 2004 and RHEL 3.
- SLAC will continue to leverage Red Hat support in enterprise environments which require it

Roadmap for RHEL 7 and CentOS 7 at SLAC

SLAC

- RHEL 7
  - Used only where Red Hat support is required, such as:
    - IBM storage servers
    - ERP business servers

- CentOS 7
  - Used everywhere else
    - Batch compute
    - Interactive login
    - Servers
    - Desktops

Chef configuration management

SLAC

- go-chef bootstrap script:
  - `$ curl yum.slac.stanford.edu/go-chef | sudo /bin/sh`
- knife commands to query chef server
  - uses /root/.chef/knife.rb file for configuration
  - `$ knife node show `hostname -f``
- Configurable attributes using: `$ knife node edit `hostname -f`
  - limit login, yum update policy, use new kernel by default or not, etc.
- system.info for Chef'd hosts (gron = grep json)
  - `/afs/slac/g/scs/systems/report/chef/system.info/`
    - `attributes and attributes.all`
    - `attributes.gron and attributes.all.gron`

Chef configuration management

SLAC

- /var/chef/cache and /var/chef/backup directories
  - Cookbooks downloaded and stored here
  - Any change made to system is recorded here
    - All previous versions of files saved here
- Chef systemctl daemon
  - systemctl status chef-client
- SLAC-CHEF github
  - Collaboration possible without the need for you to use all the Chef automate workflow
  - We add you to SLAC-CHEF github enterprise
  - You send us pull request

SLAC

Questions?

Linux Desktop Support

Karl Amrhein

June 28, 2018, Unix Town Hall Meeting

Linux Desktop Support

SLAC

- We need to define the scope of SLAC Linux Desktop Support
- What is a “Supported Linux Desktop”?
- We need to distinguish between Desktop and Server environments (not simply based on the hardware form factor)

- Desktop:
  - Mostly single-user
  - In a personal office/workspace
  - Intended for personal productivity
  - A gateway to access central Linux services
- Server:
  - Multiple users/accounts
  - Kiosk or shared workstation in a lab or control room
  - Treated as production infrastructure

Linux Desktop Support

- Target Operating Systems
  - Ubuntu 16.04 LTS (5-year Long Term Support) – 18.04 support coming soon
  - CentOS7 (10-year support, 6/2024 EOL)
- Chef Configuration Management
- Windows AD for Authentication
  - Reduce dependence on Unix Heimdal Kerberos
  - We want to converge on a single Authentication System
- Comply with SLAC MinSec (Minimum Security) policies
- Support utilities to interface with central Unix services
  - X11 Graphical terminals via FastX
  - Secure file copy and SSHFS
  - Samba (mount.cifs) for NFS
  - OpenStack Cloud
- Users can still get privs (sudo) to configure their Linux desktops
- Users will be responsible for supporting and patching their custom desktop configurations that fall outside the supported scope
- Desktop Support (SUIT) will support Linux Desktops
- SCS will continue to support Linux servers

Stanford University IT (SUIT)

SLAC

- Stanford Univerisity IT (SUIT) supports Linux desktops
  - Hands-on installation of hardware and operating system
  - Support questions and resolution

- Unix Platform
  - Maintains Chef configuration management infrastructure
  - Maintains network boot (dhcp/pxe) and automated installation infrastructure

Desktop Platforms

SLAC

- RHEL and CentOS are stable, predictable, enterprise platforms
  - Good choice for server

- Not the best desktop environment:
  - Software only gets backported – frozen at minor patch level
  - Not a polished look
  - May not be multimedia friendly out of the box
  - Getting around these issues by installing third party RPMs can quickly lead to RPM dependency issues

Ubuntu lifecycle

SLAC

- Long Term Support (LTS) releases are supported for 5 years
- Extensive software repositories
- Out of box support for displays, multimedia
- More up-to-date versions of desktop software
  - google-chrome

- Ubuntu 16.04
  - 2016 – 2021

- Ubuntu 18.04 (still need to test/verify with Chef 14)
  - 2018-2023

Kerberos Authentication

SLAC

- SLAC has two Kerberos authentication domains (realms)
  - Unix Heimdal Kerberos and Microsoft Active Directory
- Is it necessary to maintain two Kerberos realms?
- Unix is able to use Active Directory Kerberos for authentication.
- Could we one day leverage the AD Kerberos domain for all of Unix authentication? (like CERN has done)
- First step, have Linux desktops use Active Directory for authentication.
- Lots of work needed before we can do the same for servers
  - UID mapping for NFS
  - Unix host keytab generation

Kerberos Ticket Granting Ticket (TGT)

SLAC

- Active Directory TGT: user@WIN.SLAC.STANFORD.EDU
- Unix Kerberos TGT: user@SLAC.STANFORD.EDU

- You get an Active Directory TGT at login.
- You can store a Unix TGT in an alternate location and use it when needed
  - for passwordless ssh (GSSAPI)
  - Including for sshfs

```
$ env KRB5CCNAME=/tmp/krb5_cc_unix_$USER kinit -renewable $USER@SLAC.STANFORD.EDU
$ env KRB5CCNAME=/tmp/krb5_cc_unix_$USER kinit -R
$ env KRB5CCNAME=/tmp/krb5_cc_unix_$USER krenew -K60 -b -t
$ env KRB5CCNAME=/tmp/krb5_cc_unix_$USER ssh -l $USER hostname.slac.stanford.edu
```

SLAC data access using sshfs

- sshfs available in EPEL for CentOS and in Ubuntu software repository

```
$ mkdir /afs ; chown $USER /afs

$ env KRB5CCNAME=/tmp/krb5_cc_unix_$USER \
        sshfs $USER@hostname.slac.stanford.edu:/afs /afs
```

- AFS token for the sshfs mount has 25 hour lifetime
  - If the AFS space you care about has restrictive ACLs, then this matters
  - sudo umount /afs ; sshfs …
  - Use your alternate Kerberos credential cache which is renewable for 7 days

SLAC data access using Samba

SLAC

- Graphical method:
  - GNOME gvfs
  - Can browse and automatically mount via ssh, samba, etc.

- Command line method:
  - mount.cifs
    - Using command line
    - Using /etc/fstab

- Speed is comparable to NFS for a desktop on an office subnet
  - 80 MB/sec write speed

Local resources and documentation

SLAC

- Ubuntu desktop mailing list
  - Discussion list for Ubuntu on the desktop at SLAC
  - https://listserv.slac.stanford.edu/cgi-bin/wa?SUBED1=UBUNTU-L
  - ubuntu-l@listserv.slac.stanford.edu
- Chef configuration management mailing list
  - Discussion list for Chef configuration management at SLAC
  - https://listserv.slac.stanford.edu/cgi-bin/wa?SUBED1=CHEF-L
    chef-l@listserv.slac.stanford.edu
- Confluence documentation
  - https://confluence.slac.stanford.edu/display/CHEF/chef-client+daemon+mode+and+systemctl
  - https://confluence.slac.stanford.edu/display/SCSPub/CentOS+7+and+Chef

Demo

SLAC

- FastX graphical terminal - Run compute on remote server, display back to your desktop
  - Office networks are slower, and have limited connectivity as opposed to server (or farm science) networks
- sshfs demo - Create local mount point using authenticated ssh (including AFS token)
  - Can mount rhel6-64:/afs at /afs mount point on desktop - AFS tokens last 25 hours
    - Can use (alternate) local Kerberos credential cache that has renewable Unix Keberos TGT
    - For authenticated access past 25 hours: unmount, then remount (using stored Unix TGT)
    - Slower than Samba - most useful for data access that requires an AFS token,
    - Or for browsing remote filesystem - you can mount remote rhel6-64:/ and access /nfs, /afs, /scratch, etc.
- mount.cifs demo - Create local mount point using authenticated Samba
  - Can access AFS (without a token) – fine for access to non-private AFS data
  - Mostly useful for SLAC NFS access
- GNOME gvfs demo
  - Use GUI to browse remote filesystems, using Samba or SSH
  - Can create bookmarks for common access patterns
  - Mounts are done automatically (vs sshfs and mount.cifs where you specify the remote directory and the local mount point)

Questions?

CPU Hardware Refresh

Yemi Adesanya

June 28, 2018, Unix Town Hall Meeting

U.S. DEPARTMENT OF ENERGY
Office of Science

SLAC NATIONAL ACCELERATOR LABORATORY

We need to Lifecycle the Clusters

SLAC

- Majority of shared cluster nodes in B050 Datacenter are EOL or warranty has expired:
  - Bullets (Intel Westmere) >= 4 years old
  - Hequs (Intel Sandy Bridge) >= 8 years old
- Who purchased the shared cluster hardware and when it was it deployed? Go to:
  - https://confluence.slac.stanford.edu/display/SCSPub/Stakeholder+priority+on+the+Shared+Farm
- Dedicated compute silos are also EOL: SUNCAT, SIMES, etc.

New CPU footprint

SLAC

- Plans to refresh Fermi share of the clusters
- New “bubble” compute node:
  - 2 x Intel Skylake 18-core CPUs @ 2.7GHz
  - 192GB RAM or ~5.3GB per core
  - 10Gb Ethernet
  - 100Gb Infiniband
  - CentOS7 on baremetal
  - RHEL6 via containers
  - ~2500 cores in a single dense rack (*assuming power and cooling is available)
- Purchase and leasing options are available

Questions?

Singularity Containers for Batch

Renata Dart, Wei Yang

June 28, 2018, Unix Town Hall Meeting

Availability of Singularity

SLAC

Singularity is installed on CentOS7 login pool and select batch nodes

- Login pool:
  - centos7.slac.stanford.edu (load balance for cent7[a-d].slac.Stanford.edu)
- Batch nodes:
  - currently the “deft” cluster in centos7 queue
  - The “bubble” cluster (Evaluation of the Skylake CPU)
    - Fermi are benchmarking to determine if they want to buy
  - To do:
    - Change VLAN for “deft” to enable outbound TCP
    - Define batch resources: cvmfs, inet (outbound TCP), centos7, etc.
    - Singularity nodes in general fairshare queues? Requires communication…..

Singularity is also available on a few RHEL6 nodes

- For special purposes only.
- Unlikely to have everything work well under the older 2.6.x kernel.

How to Run In Singularity

Step 1: build your singularity image

- DIY: from scratch: https://singularity.lbl.gov/quickstart
  - Example singularity receipt file: /gpfs/slac/atlas/fs1/sw/slac-fermi.singularity.def
- Use prebuild images, for example, an image for Fermi & ATLAS

Step 2: run singularity on login pool or batch node

1. $ singularity exec -B /u,/scratch,/nfs,/afs,/gpfs,/cvmfs,/var /gpfs/slac/atlas/fs1/sw/slac-fermi.img.ext3 /bin/sh
   - This method is flexible and universal, can choose images and change options
   - **In addition to local file systems, we can bind mount: /afs, /gpfs, autofs (e.g. /nfs, /cvmfs, xrootdfs)**
2. $ bsub -q bubble **-app centos6** -Is /bin/sh
   - Use a pre-configurated image associate to LSF “app” **centos6**
   - **centos6** ”app” is intended to provide an environment similar (but not identical) to RHEL6 batch nodes
   - Only can only use “-app centos6” with LSF queue “centos7” and “bubble” (nodes running CentOS7)
   - When deft/bubble are part of the general fair share queue: bsub -R ”select[centos7]” -app centos6 …

Should put the prebuild images and receipts in a permanent location

Migrate Jobs from RHEL6 to Singularity Container

SLAC

ATLAS jobs: ✅

- All workflow runs - most software come from CVMFS, no dependence on AFS

Fermi jobs:

- Making good progress - thanks to Warren Focke
- **A few valuable things we learn from Fermi jobs:**
  1. Dependence on AFS
     - Require openafs rpm (for command such as 'fs')
     - @sys: Singularity inherits base CentOS7 OS's @sys (amd64_rhel70 and amd64_linux26)
       - Some of Fermi packages are under amd64_rhel60
       - Fermi choose to symlink amd64_linux26 -> amd64_rhel60
  2. OS dependent paths in PERL5LIB/@INC (similar to the @sys issue)
  3. Missing rpms, /usr/local, TWWfsw, etc.:
     - Added to the container (and the singularity receipt file)
  4. Dependent on MTA (sendmail/postfix/exim) to communicate to Fermi Pipeline
     - MTAs have setuid binaries, hard to get them working in Singularity container
     - mail -S "smtp=smtp.slac.stanford.edu:25" -s subject … works in container
     - Fermi switched to use "HTTP post" to communicate with Pipeline

Miscellaneous

Known issues:

- Singularity jobs on “deft” cluster don’t report CPU time usage correctly
  - CPU time usage freezes for running jobs
  - CPU time usage reports correctly when job finish.
  - Not seen on bubble cluster.

Other usage

- Standalone containerized applications
  - new Globus Connect CLI - aka GlobusOnline (data transfer using GridFTP)
  - Single command “globus”: many steps to install
    - installation doc assume CentOS7, not just yum, needs pip
  - Put it in a container:
    - run on both RHEL6 and CentOS7
    - $ singularity exec /afs/slac/package/globuscli/globuscli.img globus --help

Questions?

SLAC GPU Compute & Jupyter

Yee-Ting Li

June 28, 2018, Unix Town Hall Meeting

GPUs and SuperComputing

SLAC

- GPU's becoming more prevalent in Scientific Computing
  - 5 of top 7 fastest super-computers in world use GPUs
  - ORNL's Summit (1) and LLNL Sierra (3) (Nvidia Volta)
- Vector Units provides fast compute across thousands of cores per GPU
- Need smaller number of nodes
  - Sierra has 1/10th the number of cores compared to TaihuLight
- Nvidia netting $1bn income per quarter in HPC sales alone

GPU's leading computation frontier

GPU's and Scientific Computing

SLAC

- CryoEM has significant need for GPU compute
  - All/most software GPU enabled
- Typically see 10-100x increase in performance
  - heavily dependant on IO and vectorizability of code
- Increase usage of Machine Learning and Deep Learning in science
  - LCLS-2 data reduction
  - Detector signal segmentation in microDUNE
  - Event reconstruction in ATLAS

GPU's use in scientific computation becoming normal

SLAC GPU Batch Farm

- Investments from OCIO, CryoEM and HEP
  - 1 x 8 way k80
  - 1 x 8 way p100
  - 8 x 10 way 1080ti
  - 3 x 4 way v100 (32GB)
- Contact Yee if you wish to purchase GPU hardware
- Over 1 petaflops of FP32
- Access via batch queue `slacgpu`
- Will be implementing LSF fair share based on group's investment
- Open to all - submit request to unix-admin for access

SLAC GPU Batch Usage

- **Queue: slacgpu**
- **All GPUs set to ‘exclusive mode’ to prevent sharing**
- **Typically need CUDA - available via module**

```
#!/bin/bash -l
#
#BSUB -a mympi

#BSUB -W 6:00
#BSUB -q slacgpu

#BSUB -n 3
#BSUB -R "{select[ngpus>0] rusage[ngpus_excl_p=1]
span[hosts=1]}"

# enable site modules
export MODULEPATH=/afs/slac.stanford.edu/package/spack/
share/spack/modules/linux-centos7-x86_64

module load cuda-8.0.61-gcc-4.8.5-ztrhkyw
# module load cuda-9.1.85-gcc-4.9.4-fd4g2ts

<code here>
```

SLAC Jupyter Hub

- “Interactive” web based programming frontend
- Similar (same) to Google Colab
- Aimed (primarily) at Machine Learning
  - Tensorflow, Keras, Torch, NLTK, scipy, numpy etc. pre-built
- Runs on top of kubernetes (containerised)
  - Will also make use of batch system in future
- Meant for small scale programming (can fit onto one box)
- Still alpha/beta testing - Need feedback

SLAC Jupyter Hub

- **Demo**

Questions?

Future of AFS/YFS

Yemi Adesanya

June 28, 2018, Unix Town Hall Meeting

AFS at SLAC has vendor support

SLAC

- Entire SLAC AFS server infrastructure has been upgraded from OpenAFS to YFS implementation
  - Full support from vendor (Auristor)
  - 100% Linux running on x86
  - Storage volumes on SSDs
  - Significant AFS performance improvements in terms of throughput and scalability

So why are we considering the future of AFS?

SLAC

- Leading institutes (CERN, Stanford IT) are reviewing their requirement and use cases for AFS
- We have AFS dependencies in Core Unix and Science Computing infrastructure
- We should determine the scope of AFS as we move forwards (new solutions, modernization, upgrades)

- Do we carry an AFS dependency forwards to a new solution or do we eliminate it?
- Key metric: A centrally-managed HPC environment should be possible without AFS

- Configuration Management:
  - 'taylor' client management (RHEL6) requires AFS
  - Chef client management (CentOS7/RHEL7) does not require AFS

- Unix Home Directories:
  - Created in AFS by default
  - LCLS, CryoEM and ATLAS home dirs can be created in other network filesystems

- AFS protocol provides a kerberized worldwide filesystem
  - Worldwide native access was once a desirable feature for HEP collaboration
  - Native access to storage from outside the SLAC network is no longer a must-have

Steps to Reduce AFS dependencies

SLAC

- Block offsite access to AFS at SLAC
  - We ran a 24-hr test back in March 2018
  - No reports of any major impact to science workflows
  - Schedule a permanent end of offsite AFS access
- Linux Desktop “2.0” Support Model (CentOS7 and RHEL7)
  - Local home directories with Windows AD authentication
  - Access central SLAC storage and AFS via SAMBA or SSHFS
- Batch Computing
  - CentOS7 with Chef does not require AFS
  - Eliminate workflows and job submissions with hardcoded AFS links
  - Ensure batch clients are capable of running LSF without AFS
  - No immediate plans to drop AFS/YFS, but we must future-proof.

SLAC

Questions?

Virtualization as a Service (VaaS)

Ben Calvert

June 28, 2018, Unix Town Hall Meeting

Purpose

SLAC

Provide servers for customers who need to stand up applications or services

- Available quickly – hours/days instead of weeks/months
- Easy to obtain – Ticket instead of procurement
- Cheap – A fraction of the cost of physical servers
- All the benefits of virtualization with none of the headaches

Includes both science and mission support users

Hardware Purchase vs. Virtual Server as a Service

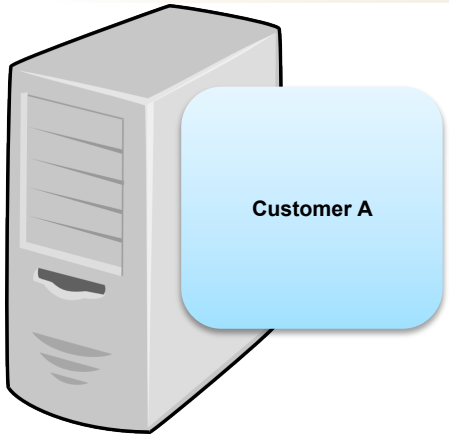

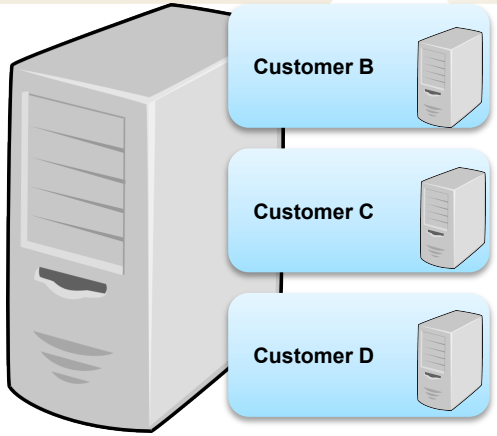

Hardware dedicated to one customer	Hardware shared among customers
Hardware is purchased (pay in full)	Service is purchased (fraction of cost)
Hardware is customer owned	Hardware is owned by OCIO
Indirect Labor provides Administration & Maintenance (Operational cost similar for both)	

Virtual servers are deployed much faster than physical servers

SLAC

	Bare Metal (Physical)	Virtual
Submit ticket	0 days	0 days
Generate quote from vendor	1 day	0 days
Purchase request, approval, submit order	5 days	0 days
Shipping, receiving, property control	10 days	0 days
Rack and stack server, power, network	10 days	0 days
Install operating system	5 days	0 days
Provision virtual machine	0 days	2 days
Cost in Time **(assumes standard configurations)**	**About 31 days**	**About 2 days**

Virtual servers allow you to buy as little capacity as you need, scale it as needed, and pay for it in small increments.

SLAC

	Bare Metal (Physical)	Virtual
Starting cost	$6,000 for server	$28,000 for hardware, software, network
Pay for needed capacity only	Not applicable	$2,000 = $28,000 / 14 capacity “blocks”
Pay over time (based on 5 year server life)	Not applicable	$400 = $2,000 / 5 years
Pay for what you need (quarterly)	Not applicable	$100 = $400 / 4 quarters
Cost in $$$	**~$6,000**	**$0 to get started** **$100 per quarter** **$400 per year** **$2,000 after 5 years**

Bare metal and virtual servers are not exact equivalents.

In some cases, virtual servers can enable you to accomplish what you need at a lower cost point.

Centralized management of hardware has benefits

SLAC

- Applications managed by Computing share the same hardware
  - Hardware used by all groups receives same level of service and support
- Computing manages the hardware lifecycle
  - Hardware is leased and will be refreshed every five years
- Centralization enables scalability and resilience
  - Centrally managed footprint has ~50 hypervisors,
    configured with vMotion for seamless maintenance of hypervisors

Virtual Server Pricing

Rates & Fees	$100 / building block / quarter (FY19) Rate evaluated annually
Minimums	Cores, memory, and storage may be ordered in “building blocks”. One building block consists of: • 2 cores • 4 GB memory • 50 GB disk **Windows** servers require **two building blocks**, at minimum Linux servers require one building block, at minimum
Term Length	No minimum term length, cancel at any time Billing occurs at the start of each quarter (not refundable)

Billing / Customer Lifecycle

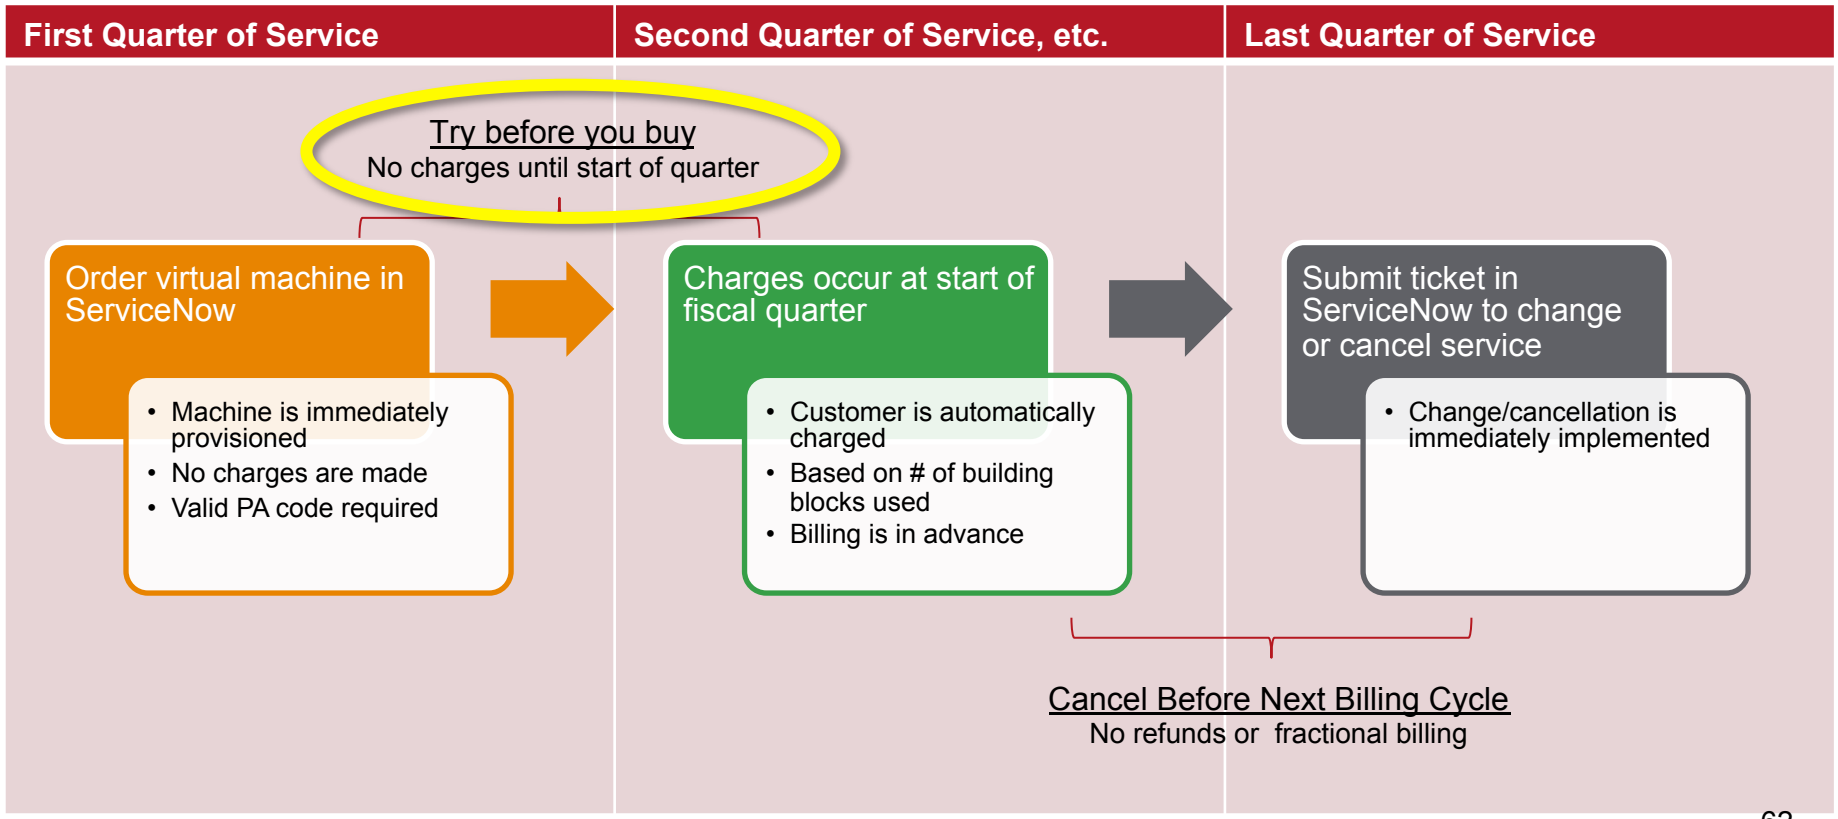

Technical details

Hardware	Dell PowerEdge R730xd/R740xd
Hypervisor	VMWare vSphere 6.5

Operating Systems	Server operating systems only
	• Linux
	  • CentOS 7 (default)
	  • RHEL 6 (upon request)
	• Windows
	  • Windows Server 2016 (default)
	  • Windows Server 2012R2 (upon request)

Networks / IP Addresses	Any data center subnet available (upon request)
System Administration (Operating System)	Fully managed by Computing Hybrid management (Computing provides patching only)
Backup Power	Two-circuit configuration with 13-minute UPS (standard) Genset power backup (by request, requires approval)
Other options	Accessible from Internet (upon request, standard process applies)

Hypothetical 1-year Service Examples (1 of 3)

SLAC

Example 1:

Testing Windows on VMaaS

1/15/2019	Request Windows Server 2016 instance with 2 building blocks	$0
1/15/2019 – 3/1/2019	Testing of server instance	$0
3/15/2019	Cancel via ticket (Testing shows bare metal server is needed)	$0
4/1/2018	Quarterly service charge	$0

No charges accrue because service is cancelled before start of next quarter

Total cost in FY 2019 : $0

Hypothetical 1-year Service Examples (2 of 3)

SLAC

Example 2:

Linux Application Server

1/15/2019	Request CentOS 7 instance with 3 building block	$0
1/15/2019 – 3/1/2019	Testing of server instance	$0
3/15/2019	Production rollout of of application on Linux virtual machine	$0
4/1/2019	Quarterly service charge	$300
7/1/2019	Quarterly service charge	$300

No charge for use of virtual machines until the start of the quarter

+ $300/quarter until service is cancelled

Hypothetical 1-year Service Examples (3 of 3)

SLAC

Example 3:

Service Modification

1/15/2019	Request CentOS 7 instance with 3 building block	$0
4/1/2019	Quarterly service charge	$300
4/15/2019	Submit request to increase building blocks to a total of 6	$0
7/1/2019	Quarterly service charge	$600

No charge for increase in building blocks until the next quarter

Total cost in FY 2019: $900

+ $600/quarter until service is cancelled

VaaS provides intrinsic benefits not included in base cloud offerings

SLAC

1. No charges for system administration of operating system
2. No data egress charges
3. No need for additional storage or security infrastructure
4. Local connectivity to storage, providing fast and seamless access to data
5. Integration with SLAC identity management
   (Linux Kerberos / Active Directory not available in AWS)
6. Compliance with Stanford and DOE cyber security requirements

Summary of VaaS Benefits

SLAC

1. Fraction of the cost of physical servers
2. Available in 1-2 days, compared to weeks for physical servers
3. Pre-built with the operating system configured
4. Easy to obtain through ServiceNow (no procurement)
5. Easy to add more memory / processors / storage (no procurement)
6. Small payments instead of one big purchase
7. Computing updates the hardware for you

How to request a virtual machine, or, a change in the Service Catalog

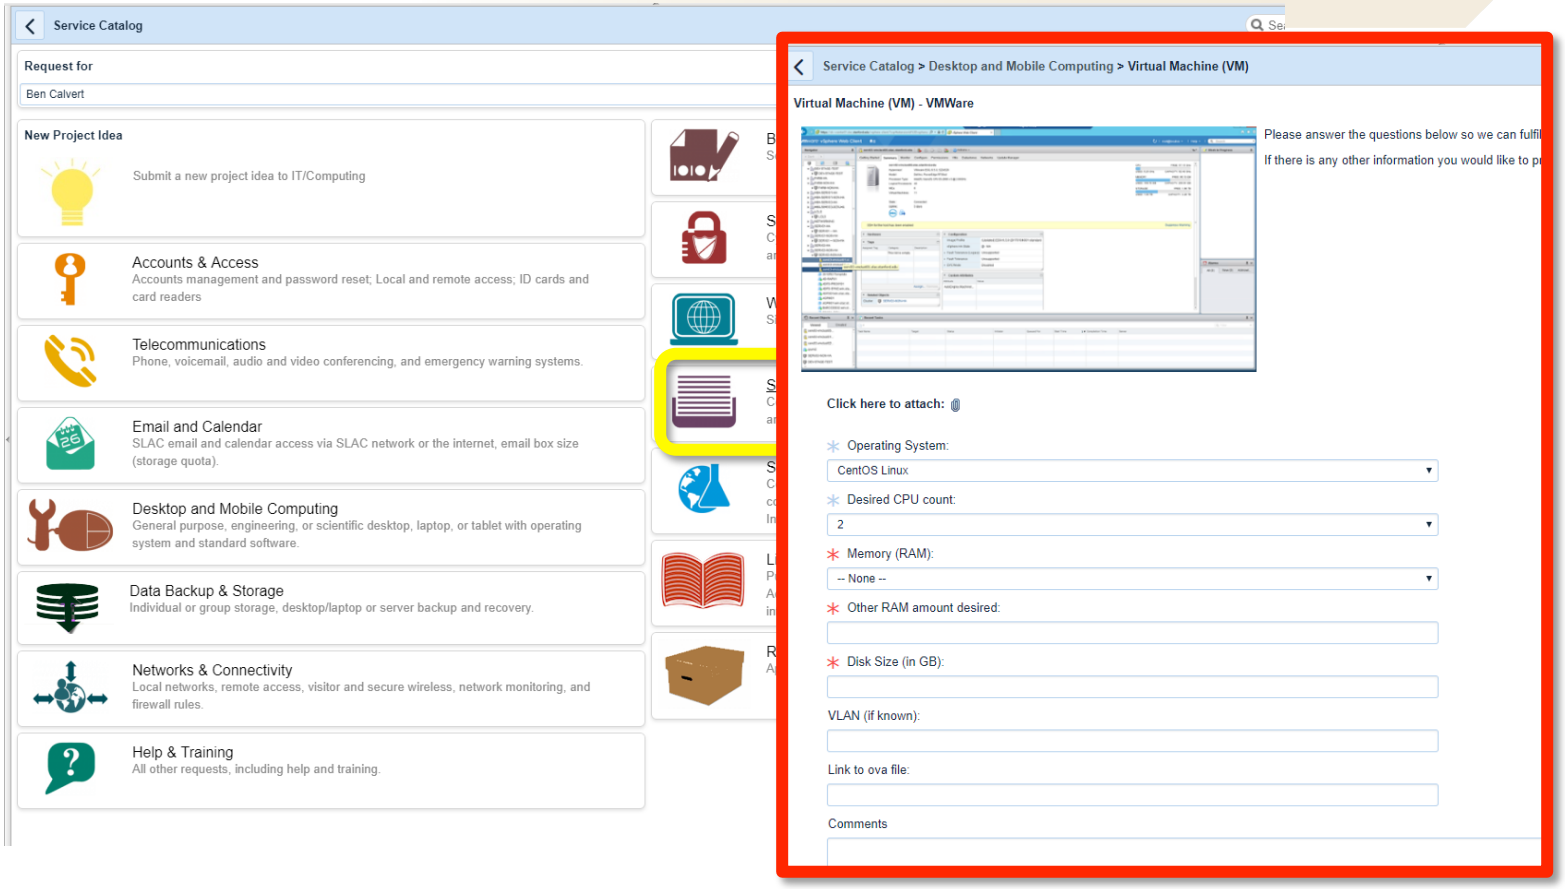

VaaS Availability

SLAC

Service becomes available October 1, 2018

No charges for virtual machines until the start of FY19

Questions?

Cyber Security Updates

Ashley Tolbert

June 28, 2018, Unix Town Hall Meeting

Agenda

SLAC

- Collaborative Accomplishments
- Audits
- On the roadmap

Collaborative Accomplishments

Web Application Scanning

- Create a process/service to scan high risk moderate applications and externally accessible applications looking for web vulnerabilities.
- Create a process to manage these vulnerabilities
- **As of May 470 out of 936 public facing external SLAC web sites scanned**
- We are at **709** Today!!!

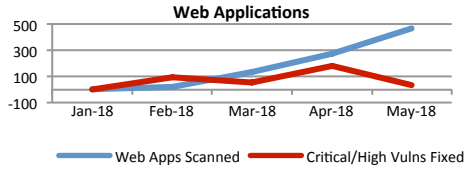

Unmanaged Systems Project

- 2015 DOE Enterprise Assessments audit finding:
  - Insecure and potentially compromised devices can connect to SLAC Network
- VPN and Visitor wireless network are used for SLAC business on unmanaged devices (SLAC-owned or employee-owned)
- **Completed**: VPN Posture assessment
- **Up next**: Visitor Wireless

Unmanaged Systems

SLAC

Problem

- 2015 DOE Enterprise Assessments audit finding:
  - Insecure and potentially compromised devices can connect to SLAC Network
- US Government FISMA targets
  - All devices have a posture assessment before connecting to VPN
- Stanford University
  - All devices must comply with minimum security standards before connecting
- VPN and Visitor wireless network are used for SLAC business on unmanaged devices (SLAC-owned or employee-owned)
- **Visitor wireless is a mission-critical capability used to deliver services to facility users**

Objective

- Reduce the risk posed by unmanaged systems

Approach

- **Completed:** VPN (March 19th – Blocking enabled)
  - Enable Cisco ASA feature to interrogate systems attempting to get on the VPN

Consistent, steady progress in resolving audit/assessment findings…

Year	Auditor	Title	Observations	Status
External Audits/Assessments				
2013	DOE Inspector General	FISMA	9 findings	All closed
2015	DOE Enterprise Assessments	Cyber Review	7 findings, 11 observations	1 open finding
2016	DOE Inspector General	FISMA	None	N/A
2018	DOE Office of Science	Safeguards & Security	0 findings, 2 observations	N/A
2018	DOE Inspector General	Legacy infrastructure	April 2018	April 2018
2018	DOE Inspector General	Cyber and IT Controls	May 2018	May 2018
Internal Audits/Assessments				
2014	Stanford Internal Audit	PeopleSoft Readiness	5 findings	All complete
2016	SLAC/Stanford Internal Audit	A-123	None	N/A
2017	SLAC	A-123	None	N/A

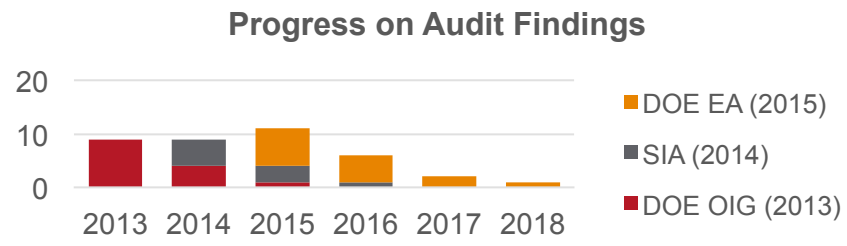

…and in avoiding new findings by minding our house.

On The Roadmap

- Community outreach for Visitor wireless project

- Cyber Documentation Updates
  - Science enclave System Security Plans
  - Business Impact Assessment
  - Threat Assessment

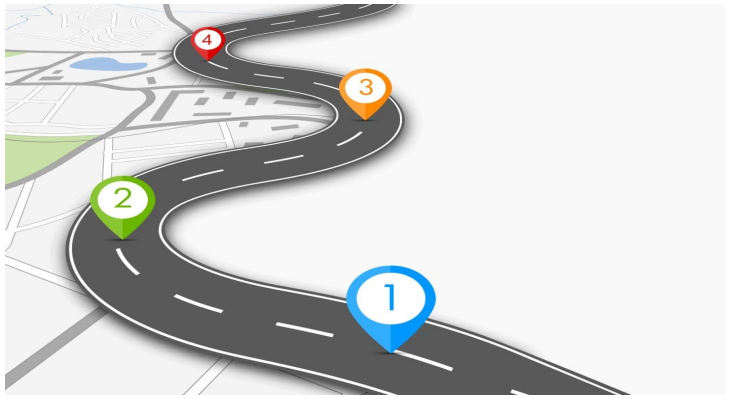

Questions?