Computing Division
Unix Town Hall

Yemi Adesanya, Associate Director Scientific Computing Services

November 14th 2019

U.S. DEPARTMENT OF ENERGY
Office of Science

SLAC NATIONAL ACCELERATOR LABORATORY

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

Scientific Computing Services

Scientific Computing Services (confluence) page

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

unix-admin@slac.stanford.edu

support/questions

yemi@slac.stanford.edu

650-926-2863

Unix Town Hall Meeting

Agenda:

- Conferences & Training
- SDF Strategy
- SDF Phase 1 Deployment
- SDF Architecture (Shawfeng)
- SLAC Networking for SDF (Mark F.)
- Questions / Break
- NERSC (Debbie)
- CentOS / RHEL Platform Update (Karl)
- Storage & Data Management (Lance)
- Cyber Security Update (Cyber)
- Questions/Discussion

Conferences and Training

- **SC19**, November 17 - 22, Denver, CO, USA
  - https://sc19.supercomputing.org
- **KubeCon 2019**, November 19 - 21, San Diego, CA
  - https://events19.linuxfoundation.org/events/kubecon-cloudnativecon-north-america-2019/
- **NeurIPS 2019**, December 8 -14, Vancouver, Canada
  - https://nips.cc/
- **Internet2 TechEx 2019**, December 9 - 12, New Orleans, LA
  - https://meetings.internet2.edu/2019-technology-exchange/
- **GTC 2020**, March 22 - 26, San Jose, CA
  - https://www.nvidia.com/en-us/gtc/

Shared Science Data Facility (SDF) Strategy

Yemi Adesanya

November 14th 2019, Unix Town Hall

U.S. DEPARTMENT OF ENERGY
Office of Science

SLAC NATIONAL ACCELERATOR LABORATORY

SDF Goals

- Deliver **common shared computing infrastructure** to tackle massive throughput data analytics at SLAC

- Enable critical, data-heavy computing workflows in several key mission areas:
  - **SLAC users facilities (LCLS, UED, CryoEM, SSRL), Machine Learning, HEP, FES**

- The SDF infrastructure would offer:
  - **High-throughput** and **high capacity storage**
  - Comprehensive set of frameworks, tools and services
  - **Baseline capabilities for all SLAC users**
  - A cost model for stakeholders with demands that exceed the baseline

SDF Goals

- The benefits of a centrally integrated hardware architecture:
  - **Increased operational efficiencies** (lower administration overhead)
  - Coordinated procurements for Economies of Scale
  - Increased utilization by leveraging ‘idle/free’ compute cycles

- Promote a model for **sustainable scientific computing services**
  - Drive lifecycle and continued support for modern, capable solutions to deliver the science

- **Strong Alignment to Science Goals and Priorities**
  - Partner with science via SDF steering and advisory committees

The Challenge: Raise the Bar for “Baseline” Scientific Computing

- What do we consider to be “Baseline”?
  - A Service that addresses a **common computing requirement** (not unique to any specific project or application)
  - A Service typically **managed and supported through a central organization** (OCIO SCS team)
  - A level of **service the user community expects from the lab as a “birthright” entitlement (for free)**

- Why is Baseline Scientific Computing important?
  - **Baseline services are at the core of many (critical) Scientific applications**
  - **Baseline capabilities help seed new science initiatives before any project-specific grants are awarded**
  - **Baseline solutions foster labwide collaboration and partnership**

The Challenge: Raise the Bar for “Baseline” Scientific Computing

- What is the risk posed by lack of support for Baseline?
  - **No ongoing strategy to address the current and future core computing needs of the lab**
  - **No sustainable lifecycle or modernization**
  - **Decentralization leads to inefficiencies, lack of governance, policy, etc**
  - **Science and collaboration suffers**

Making the Case for Shared Integrated Infrastructure

- **LCLS-II and CryoEM** applications/workflows demand similar high-throughput solutions
- **LCLS-II infrastructure could potentially contribute to the Baseline Capability**
    - 70% of LCLS-II compute time could run other science without impacting LCLS-II operations
- **SLAC Machine Learning initiative** also requires optimal bandwidth between compute (GPU) and storage
- Integrate compute and storage hardware projections from these facilities/projects
- Architect a common infrastructure and consider scalability and total operating cost

	Stage 1 (2019-2024)	Stage 2 (2025-2028)	Main Driver
CPU Compute	1 PFLOPS		LCLS-II
GPU	1 to 10 PFLOPS	> 10 PFLOPS	Cryo-EM + LCLS-II + ML
Disk Storage	10 to 30 PB	50 to 100 PB	Cryo-EM + LCLS-II
Tape Archive	10 to 100 PB	100 to 500 PB	HEP + LCLS-II
Border network	200 Gb/s	1 Tb/s	LCLS-II

SDF is NOT about deploying Siloed Solutions

SLAC

- Our **existing siloed solutions**:
  - Are **Inefficient** in terms of scalability, utilization and support
  - **Hinder sustainability** of compute, network and storage resources
  - Prevent implementation of **baseline services** to provide meaningful resources for all users; complicates use
  - Impact **long-term planning**

Silos limit our ability to collaborate and align on Computing Strategy!

So what exactly is SDF?

SDF is more than a "facility", it's an overarching Computing Strategy

- **An integrated hardware design that includes Storage, Compute, GPU and Fast Networking?**
  - Yes, all of the above. The focus is on fast access to storage from the compute servers
- **A funding model for all of this hardware?**
  - Yes, SDF will standardize on limited number of hardware configurations and coordinate combined purchases with stakeholders / business managers
- **A Datacenter Strategy?**
  - Yes (See Christian Pama). We need to carefully plan for the future infrastructure as it scales over time
- **An Organization?**
  - We'll develop a matrixed organization of talent distributed across the lab. It will take an entire village to pull this off!
  - SDF will be overseen by a steering committee comprised of key science representatives to ensure alignment with Mission requirements and priorities
- **A set of policies and best practices?**
  - SDF must ensure resources are managed effectively through policies and controls
  - (examples: storage quotas, hardware lifecycle refresh, data retention periods)
- **Raise the bar for Baseline Scientific Computing**
  - Seek lab funding to sustain the baseline
  - Share project resources (LCLS-II) when feasible
  - Continual engagement with the Science Community to stay aligned with evolving requirements

SDF Phase 1 Deployment

Yemi Adesanya

November 14th 2019, Unix Town Hall

SDF Phase 1 Storage

Phase 1 Storage is onsite and we are preparing for installation

- Two DDN ‘Exascaler 5.0’ Lustre 18K controllers for LCLS-II, CryoEM, and Baseline (lab-funded)
  - Up to ~70GB/sec per controller
  - Up to 1800 HDDs per controller
- ~250TB SSD pool for home dirs
- ~7PB HD ‘data’ storage
- Declustered RAID, Distributed metadata
- Expand a single namespace across multiple controllers and storage pools
- Future Option: Automatic tiered storage between NVMe and Disk

What about Storage-as-a-Service?

- StaaS will be replaced with SDF Baseline storage in 2020
- SDF Baseline is lab-funded! But "free" only up to a certain point!
- We'll develop some initial user and project quota limits and usage guidelines
- Expect Baseline will cover modest project requirements (50-100TB)
- Bigger demands (several PBs) will need project funding
- SDF storage is optimized for SDF compute
- We will limit access from legacy environments
- The intent is to build up SDF as we retire older clusters

SDF Phase 1 GPU

Integrating 11 new Baseline funded GPU nodes as part of SLAC Machine Learning initiative. Thank you, Daniel Ratner!

- Dual Intel Skylake 12-Core Processors
- 192GB RAM
- 10 x 2080Ti (11GB Mem)
- 6TB local SSD “scratch”

Existing CryoEM GPU servers will also be migrated to SDF

SDF Phase 1 CPU

Combined purchase of CPU cores funded by LCLS-II, Fermi (HEP) and SUNCAT (BES)

- Procure initial ~0.2 PFLOP by January 2020
- AMD Rome or Intel Cascade Lake depending on price and code-specific performance
- Minimum of 4GB RAM/core ratio
- 100Gb Infiniband for high-throughput to SDF storage
- 10Gb ethernet
- Flexible cluster management and provisioning to support priority workloads

SDF is a greenfield for modern solutions

SDF migration will not be seamless, but we need to modernize

- Slurm is currently our preferred batch scheduler for SDF
  - Comprehensive support for GPU scheduling (fairshare)
  - Widely used within the research computing community (SRCC, NERSC, etc)
- Active Directory authentication
  - Integrate with open-source Identity Management framework
  - Avoid building dependencies on dated, homegrown infrastructure
  - More potential for streamlined and automated account provisioning
- CentOS 7
  - Run legacy RHEL6 applications in Singularity containers

Stakeholder requirements will shape SDF capabilities

- The success of SDF will be measured on how we align with the science needs
- SDF stakeholders will provide requirements through a steering committee
- Hardware will be purchased and lifecycled periodically based on current supported standards
- We anticipate a heterogenous (but controlled) hardware environment as technology and requirements evolve
- We must be flexible and responsive

Questions?

SDF Architecture

Shawfeng Dong

November 14, 2019, Unix Town Hall

U.S. DEPARTMENT OF ENERGY
Office of Science

SLAC NATIONAL ACCELERATOR LABORATORY

CPU Nodes

- Dell PowerEdge C6525 Servers

- Node Specs:
  - 2x 64-core AMD *Rome* EPYC 7702 CPUs @ 2.0GHz
    - AVX-256 SIMD
  - 512GB RAM (4GB per core)
  - Mellanox ConnectX-6 100Gb/s HDR100 InfiniBand Adapter
  - 10GbE Base-T Ethernet
  - 960GB SSD

Compared to Bubble Nodes

- Dell PowerEdge C6420 Chassis
- Node Specs:
  - 2x 18-core Intel *Skylake* Xeon Gold 6150 CPUs @ 2.7GHz
    - AVX-512 SIMD
  - 192GB RAM (~5.3GB per core)
  - Mellanox ConnectX-4 EDR 100Gb/s InfiniBand Adapter
  - 10GbE SFP+ Ethernet
  - 960GB SSD

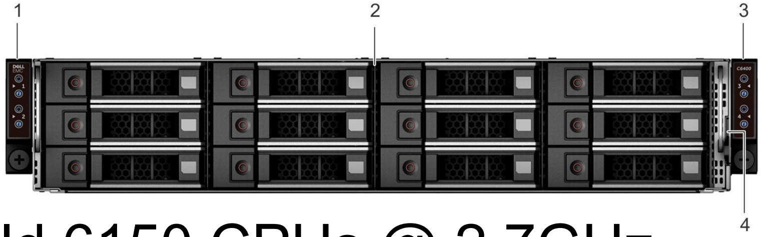

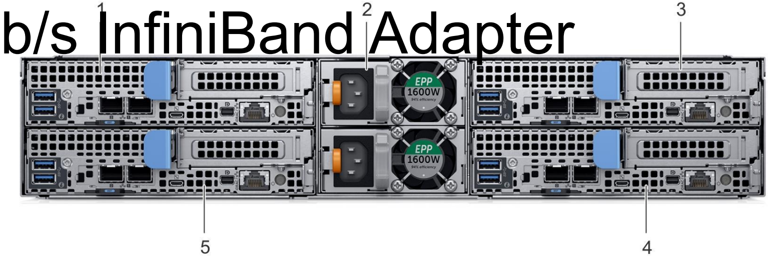

GPU Nodes

SLAC

- Supermicro SYS-4029GP-TRT2
- Node Specs
  - 2x 12-core Intel *Skylake* Xeon Gold 5118 CPUs @ 2.3GHz
  - 192GB RAM
  - 10x Nvidia GeForce RTX 2080 Ti
    - 4352 CUDA cores @ 1.35GHz
  - 10GbE Base-T Ethernet
  - 1x 480GB + 3x 1.92TB SSDs

Fat-Tree Topology (2:1 oversubscription)

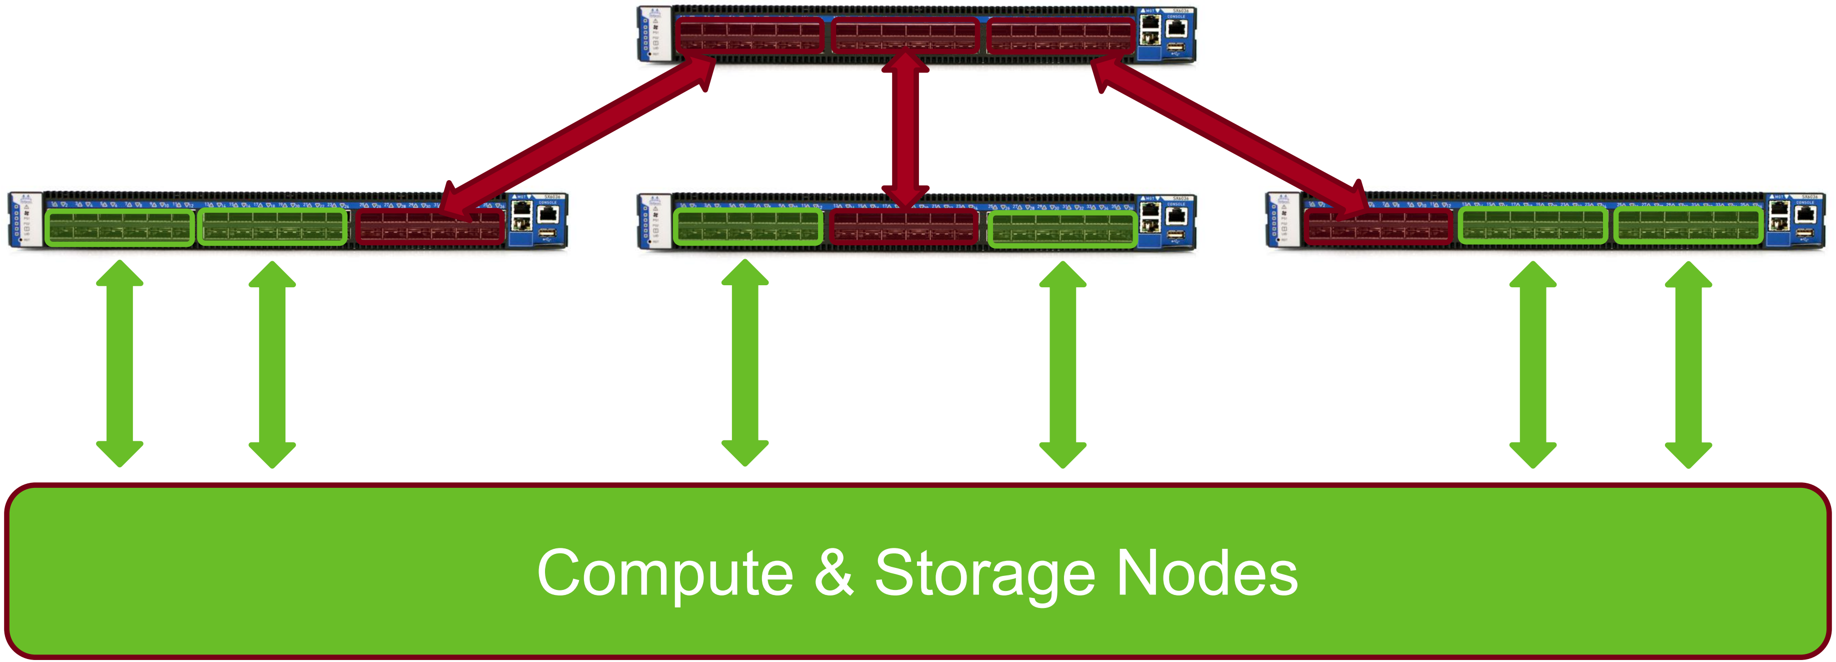

Future Expansion

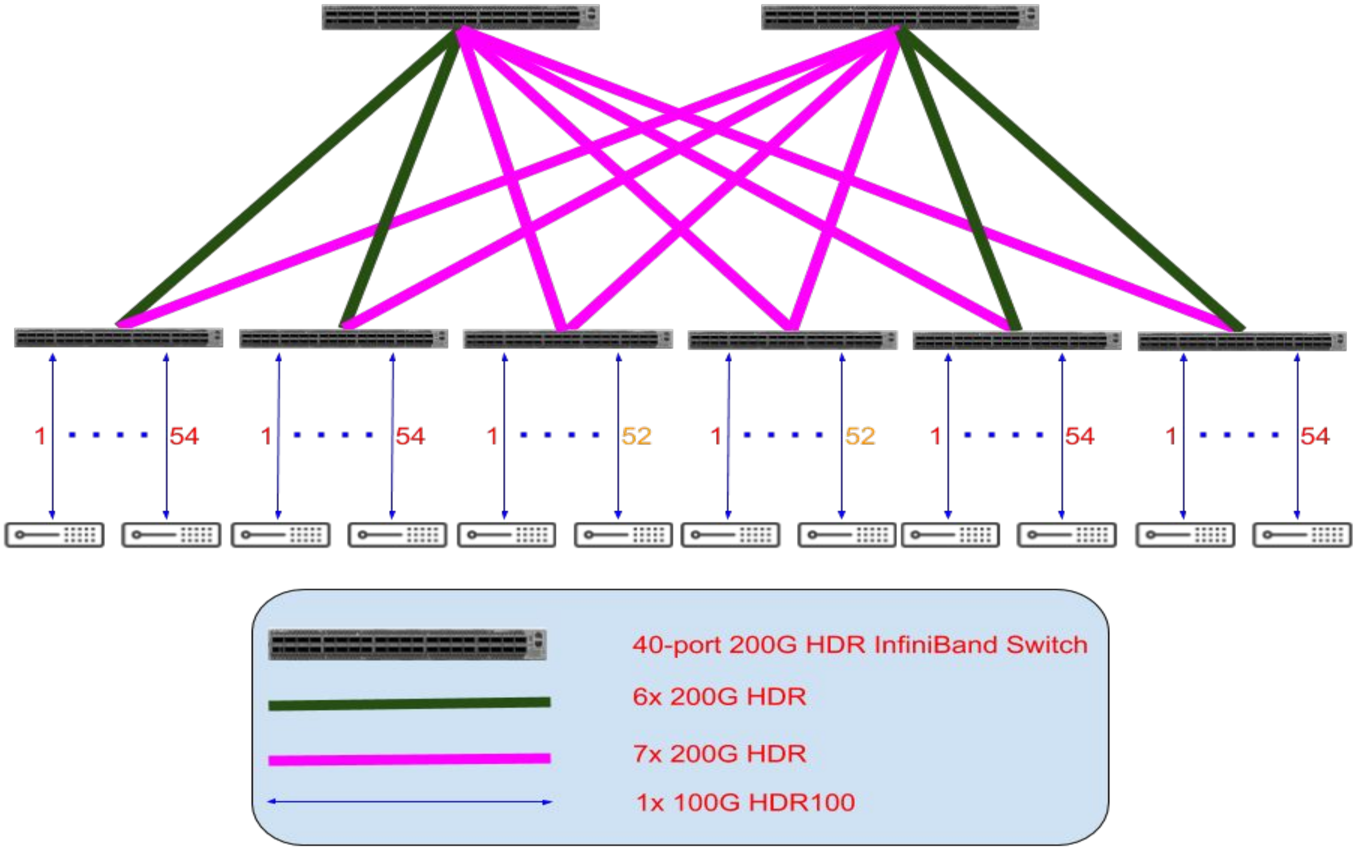

Preliminary Rack Elevation

SLAC

Water-Cooled Racks

U	1BD22	1BD21	1BD19	1BD17	1BD13
47	10G Ethernet	10G Ethernet	10G Ethernet	10G Ethernet	
46					
45	IB Leaf 1	IB Spine	IB Leaf 2	IB Leaf 3	
44					
43	4x CPU Nodes	4x CPU Nodes	4x CPU Nodes	4x CPU Nodes	
42					
41	4x CPU Nodes	4x CPU Nodes	4x CPU Nodes	4x CPU Nodes	
40					
39	4x CPU Nodes	4x CPU Nodes	4x CPU Nodes	4x CPU Nodes	
38					
37	4x CPU Nodes	4x CPU Nodes	4x CPU Nodes	4x CPU Nodes	
36					5x 4.333U GPU Nodes
35	4x CPU Nodes	4x CPU Nodes	4x CPU Nodes	4x CPU Nodes	
34					
33					
32	Login Node	Login Node	Login Node	Login Node	
31					
30	1U GPU Node	1U GPU Node	1U GPU Node	1U GPU Node	
29	1U GPU Node	1U GPU Node	1U GPU Node	1U GPU Node	
28					
27					
26					
25					
24					
23					
22					
21					
20					
19					
18					
17					
16					
15					
14	5x 4.333U GPU Nodes	5x 4.333U GPU Nodes	5x 4.333U GPU Nodes	5x 4.333U GPU Nodes	5x 4.333U GPU Nodes
13					
12					
11					
10					
9					
8					
7					
6					
5					
4					
3					
2					
1	OOB	OOB	OOB	OOB	OOB

Computing Environment

- *New* SDF Batch Scheduler: SLURM
- We use the *module* utility to manage software environment for HPC
  - Compilers
  - MPI
  - Libraries
  - Machine Learning Frameworks
- Singularity container engine

Questions?

SLAC Networking for SDF

Mark Foster

November 14, 2019, Unix Town Hall

U.S. DEPARTMENT OF ENERGY
Office of Science

SLAC NATIONAL ACCELERATOR LABORATORY

SLAC Networking for SDF - 2020

Initially: support 200 Gbps (aggregate) capability between SLAC and other sites using multiple ESnet 100G links; ultimately scale to Nx400 Gbps

Connection speeds subject to change as newer technology matures

Resiliency: possible reduced bandwidth on any single link failure

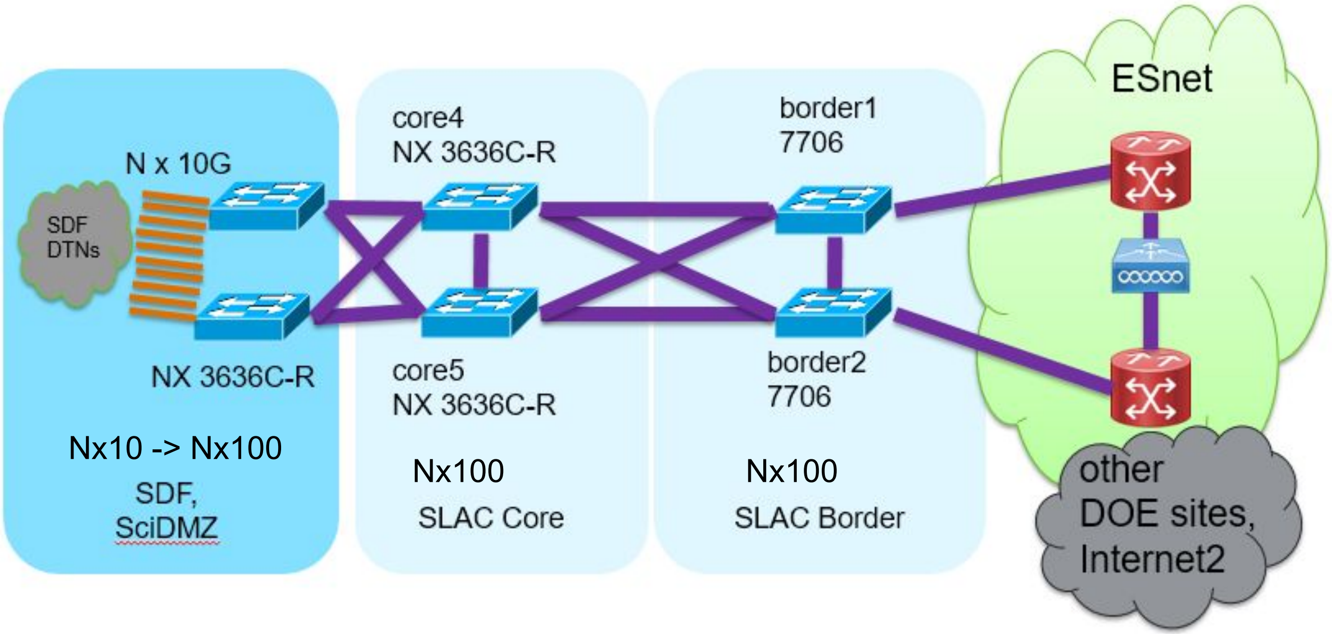

Ability to scale Nx100G (now) and Nx400G (future)

SLAC/ESnet Projections

- ESnet: primary WAN provider for SLAC
- SLAC capabilities keep up until 2023-2024
- Technology increases for 1 Tbps connections still in development
- Interim solutions by Nx100G do not scale with logarithmic growth in usage
- Projected plans:
- FY23: upgrade existing core + border
- FY25: new core, border

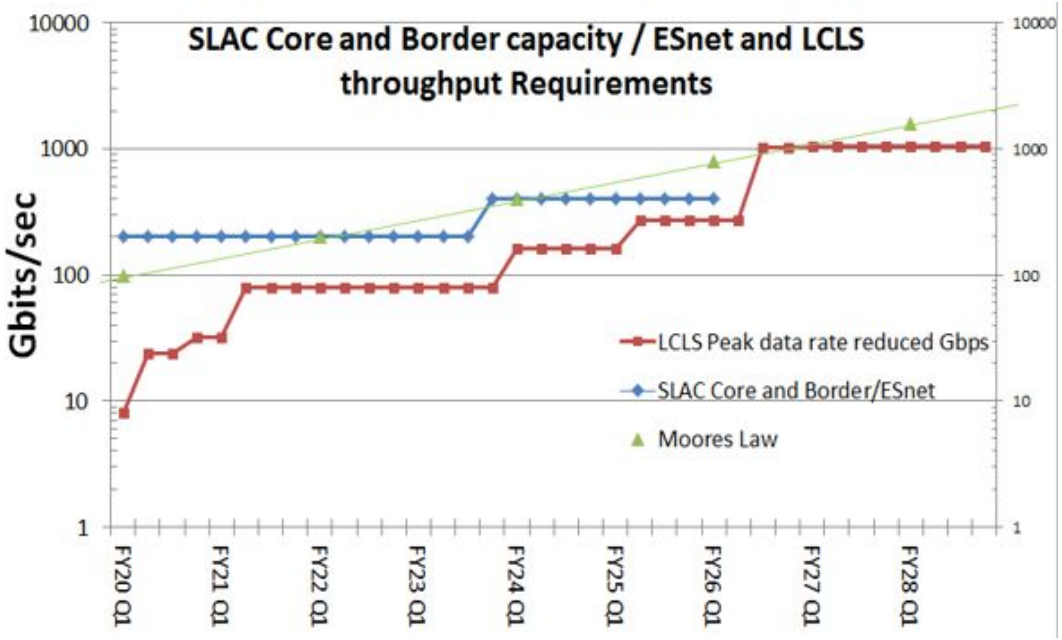

Current forecast feasible, demands from combined projects may push faster upgrade

5min Intermission / Questions?

NERSC Updates

Debbie Bard

November 14th 2019, Unix Town Hall

U.S. DEPARTMENT OF ENERGY
Office of Science

SLAC NATIONAL ACCELERATOR LABORATORY

CentOS / RHEL Platform Update

Karl Amrhein

November 14th 2019, Unix Town Hall

CentOS and RHEL server lifecycle dates

OS Distribution	General Availability	Vendor Retirement	Extended Lifecycle Support
RHEL 5	2007-03	2017-03	2020-11
RHEL 6	2010-11	2020-11	2024-06
RHEL 7	2014-06	2024-06	TBD
CentOS 7	2014-07	2024-06	None
RHEL 8	2019-05	2029-05	TBD
CentOS 8	2019-09	2029-05	None

10 years of vendor support, plus Extended Lifecycle Support (RHEL only)

CentOS and RHEL server plans and priorities

SLAC

- Priority: any remaining RHEL 5 servers need to be retired or migrated to CentOS 7
  - CentOS 7 is the recommended distribution for centrally managed linux servers
  - Choose RHEL when it appropriate to pay for vendor support for that server
  - We have thousands of batch compute nodes, and hundreds of dev/test nodes

- CentOS started in 2004 as an independent distribution
  - Since 2014 CentOS has been sponsored by Red Hat
  - Most CentOS project lead developers are Red Hat employees

- Ubuntu 18.04 LTS is the recommended linux distro for Desktops (or CentOS 7.x)
  - We do not have the staffing resources to centrally support Ubuntu on server platforms (eg, see previous matrix which has 6 linux distros for servers)

Chef Configuration Management update

- We continue to maintain two configuration management systems:
  - Taylor is used on RHEL 6 or earlier; Chef is used on CentOS / RHEL 7 or later
- Chef Infra Server - running version 12; upgrade to 13 needed before April 2020
- Chef Infra Client - running version 14; upgrade to 15 needed before April 2020
- Chef Automate (operational dashboard) - running version 2.latest, updated automatically
- Chef vendor and community support:
  - two slack workspaces: Chef Success, and Chef Community
  - Account Manager (Denise) and Solutions Architect (Jeff)
  - Ticketing system for support: https://getchef.zendesk.com
  - When possible, we use community/vendor cookbooks via “wrapper” cookbooks

Chef Source Code Management, Testing, and Delivery

- https://github.com/SLAC-CHEF/
  - email unix-admin if you wish to get access; collaboration and PRs welcome!
- We have a Jenkins Automated Testing and Delivery Pipeline
  - each chef cookbook is a github repo in the SLAC-CHEF organization
  - the build/test/deploy pipeline is codified and stored in each repo (Jenkinsfile)
  - A “git push” initiates the automated testing and delivery pipeline (CLI or GitHub)

  -> github sends a webhook to a slack channel, and to jenkins

  -> jenkins does a git checkout of updated code

  -> jenkins spins up new VM(s), and chef code applied (vagrant)

  -> jenkins does lint, unit, and integration testing (test kitchen)

  -> jenkins sends a slack webhook for human approval

  -> jenkins delivers code to chef server

Chef logs, reports, and node/role configuration

- As root on your server:
  - journalctl -u chef-client
  - knife node show `hostname -f` -F json
  - /var/chef/backup/ { etc/ root/ usr/ var/ }
  - /var/chef/cache/cookbooks/
- Node attributes and chef roles saved in AFS git repositories
  - /afs/slac/g/scs/systems/report/chef/ { system.info/ roles/ }
- Nodes can have individual or role attributes (options) for configuration
  - some familiar configuration options you have seen in /etc/taylor.opts for RHEL 6
  - sudo privs, list of users/groups who can log in, increment default kernel or not, etc
  - groups of nodes can get identical configuration with one line using roles
    - some examples: batch, spear, lsst, fileserver, casper, etc.

Questions?

Storage & Data Management

Scientific Computing Services

Lance Nakata, November 14, 2019

U.S. DEPARTMENT OF ENERGY
Office of Science

SLAC NATIONAL ACCELERATOR LABORATORY

Storage Update

- Storage as a Service (StaaS)
  - StaaS has been under a heavy load lately. Slowdowns are noticeable.
  - Some changes and statistics:
    - Increased disk count from 180 (714TB) to 240 (874TB) to help out
    - Will still need to rebalance disks for smoother I/O performance
    - Quota allocated: 968TB (683TB in Feb)
    - Quota in use: 612TB (393TB in Feb)
    - Space available: 262TB (321TB in Feb)
    - Files on disk: 286M (263M in Feb)
  - We will be installing new storage and starting migrations thereafter
  - BTS Director Suzanne Hansen has declared current StaaS GPFS allocations to be free for FY20 as we lifecycle to new storage

Storage Update (2)

- Tape
  - Looking at better ways to integrate with disk to reduce storage costs
  - Trying not to buy more T10K media due to obsolescence
  - New silo in our future, probably with LTO tape drives
- Hardware Lifecycle
  - Still actively retiring Sun and LSI hardware
  - Some migrations are dependent upon new SDF storage deployment
  - You will receive migration (downtime) notices for data moves

Questions?

Cyber Security Update

Michelle Jost, SLAC Cyber Security Team

November 14th, 2019

U.S. DEPARTMENT OF ENERGY
Office of Science

SLAC NATIONAL ACCELERATOR LABORATORY