Computing Division

Scientific Computing Services

Unix Town Hall Meeting

Yemi Adesanya, March 2, 2017

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

Agenda:

- Announcements
- IS3C
- SLAC<->NERSC Partnership
- NERSC Overview
- *5 min break*
- Storage & Data Management
- Unix Platform
- Container Technology
- Questions/Discussion

Announcements

Scientific Computing Services

Yemi Adesanya, March 2, 2017

Conferences and Training

- NVIDIA GPU Technology Conference
  - May 8th-11th in Silicon Valley
- Red Hat Summit
  - May 2nd-4th Boston, MA
- ChefConf 2017
  - May 22nd-25th Austin, TX

IS3C

Integrated & Sustainable SLAC Scientific Computing

Yemi Adesanya, March 2, 2017

How the Scope for IS3C Emerged

- OCIO performed aging hardware analysis on science systems
- Recognized science hardware at risk
  - ~75% of Scientific Compute cores are > 5 years old and have no hardware warranty
  - > 6PB of Science Data on storage > 5 years old with no hardware warranty
- Began socializing aging risk of systems to science community
- Additional context emerged regarding other science needs
- These needs were holistic in that they tied together people, process, & technology (workforce gaps, inadequate policies and funding, disparate systems, and non-integrated science requirements)
- IS3C scope emerged through socialization
- Enterprise-level risk added to Lab Risk Registry owned by the CIO – We Start From Here

Recurring Themes/Challenges Based on Socialization

SLAC

- One-time program-based funding has led to aging and inadequate compute and storage infrastructure to support Lab Objectives and Agenda
- Unsustainable processes to support gathering science requirements to determine optimal facilities footprint
- Lack of holistic approach to integrated & optimized scientific computing services: policies, tools, workforce planning, sustainable funding models
- Lack collaborative Leadership approach
- Growing data management needs and concerns

IS3C Socialization with SLAC Science Council & Senior Leadership

- Science Community (Mission)
  - SLAC Director and Deputy Director
  - David MacFarlane
  - Tom Abel *(Planned)*
  - Alex Aiken
  - Phil Bucksbaum *(Planned)*
  - Mike Dunne
  - Mike Fazio
  - Mark Hartney
  - Tony Heinz
  - JoAnne Hewett
  - Keith Hodgson (*Planned*)
  - Kelly Gaffney
  - John Galayda
  - Siegried Glenzer / Frederico Fuiza
  - Steve Kahn
  - SRCC (Ruth Marinshaw)
  - Lia Merminga
  - Despina Milathianaki *(Planned)*
  - Richard Mount
  - Jens Norskov
  - Aaron Roodman *(Planned)*
  - Robert Schoenlein *(Planned)*
  - John Seeman
  - Z-X Shen
  - Soichi Wakatsuki
  - Bill White

- Mission Support
  - Marc Clay (Contractor Assurance)
  - Charlotte Chang
  - Paul Chiames (HR)
  - Suzanne Davidson (CFO)
  - Susan Simpkins (Business Technology Services)
  - Steve Nott (Procurement)
  - Russ Thackston and Jeff Sims (Facilities)

- DOE Site Office
- IT Independent Review Board
- SLAC OCFO
- DOE OCIO

IS3C Advisors:

- Richard Mount
- Frederico Fuiza
- Johannes Voss
- Henry van den Bedem
- Tony Johnson

IS3C Addressing Current Challenges in SLAC Scientific Computing

SLAC

Area	Current State/Challenges	How IS3C Addresses
Requirements	Piecemeal, program-based gathering of requirements as needed	*Project Track: Requirement* - Streamlined and sustainable requirements gathering framework with integration where applicable
Tools and Algorithms	Unclear understanding of critical tools used by the scientific community. Limited grouping of purchases leading to non-optimal licensing costs	*Project Track: Tools and Algorithms* - Scientific Computing Toolkit (part of large set of engineering, administrative and collaboration toolkits) - Licensing cost assessment and recommendation
Compute, Storage, Network	Aging unfunded infrastructure inadequate to meet future needs. Infrastructure scattered throughout Lab with unclear understanding of use	*Project Track: Compute, Storage, Network* - Understanding of computing footprints across the Lab - Lifecycled Infrastructure aligned with requirements - Efficient and sustainable process to review new technologies to meet requirements - Dashboards and metrics measuring sustained operations
Scientific Computing Workforce	Lab workforce not optimized to meet Lab needs and unable to leverage critical skillsets across different science programs	*Project Track: Resource Capabilities* - Clearer understanding of roles and responsibilities, critical dependencies, training - Current and future state workforce, including gaps - Recruitment and Retention strategy
Policies	Many missing policies leading to risk-based ongoing practices (e.g. lack of a data management policy)	*Project Track: Policies* - Clear documented policies aligned with process and funding (e.g. hardware lifecycle, data management, software management, etc.)
Funding Models	Unsustainable to support service-based models, address hardware lifecycle, and future needs	*Project Track: Sustainable Funding Models* - Develop sustainable funding models to support Lab infrastructure needs and services

Current Requirements

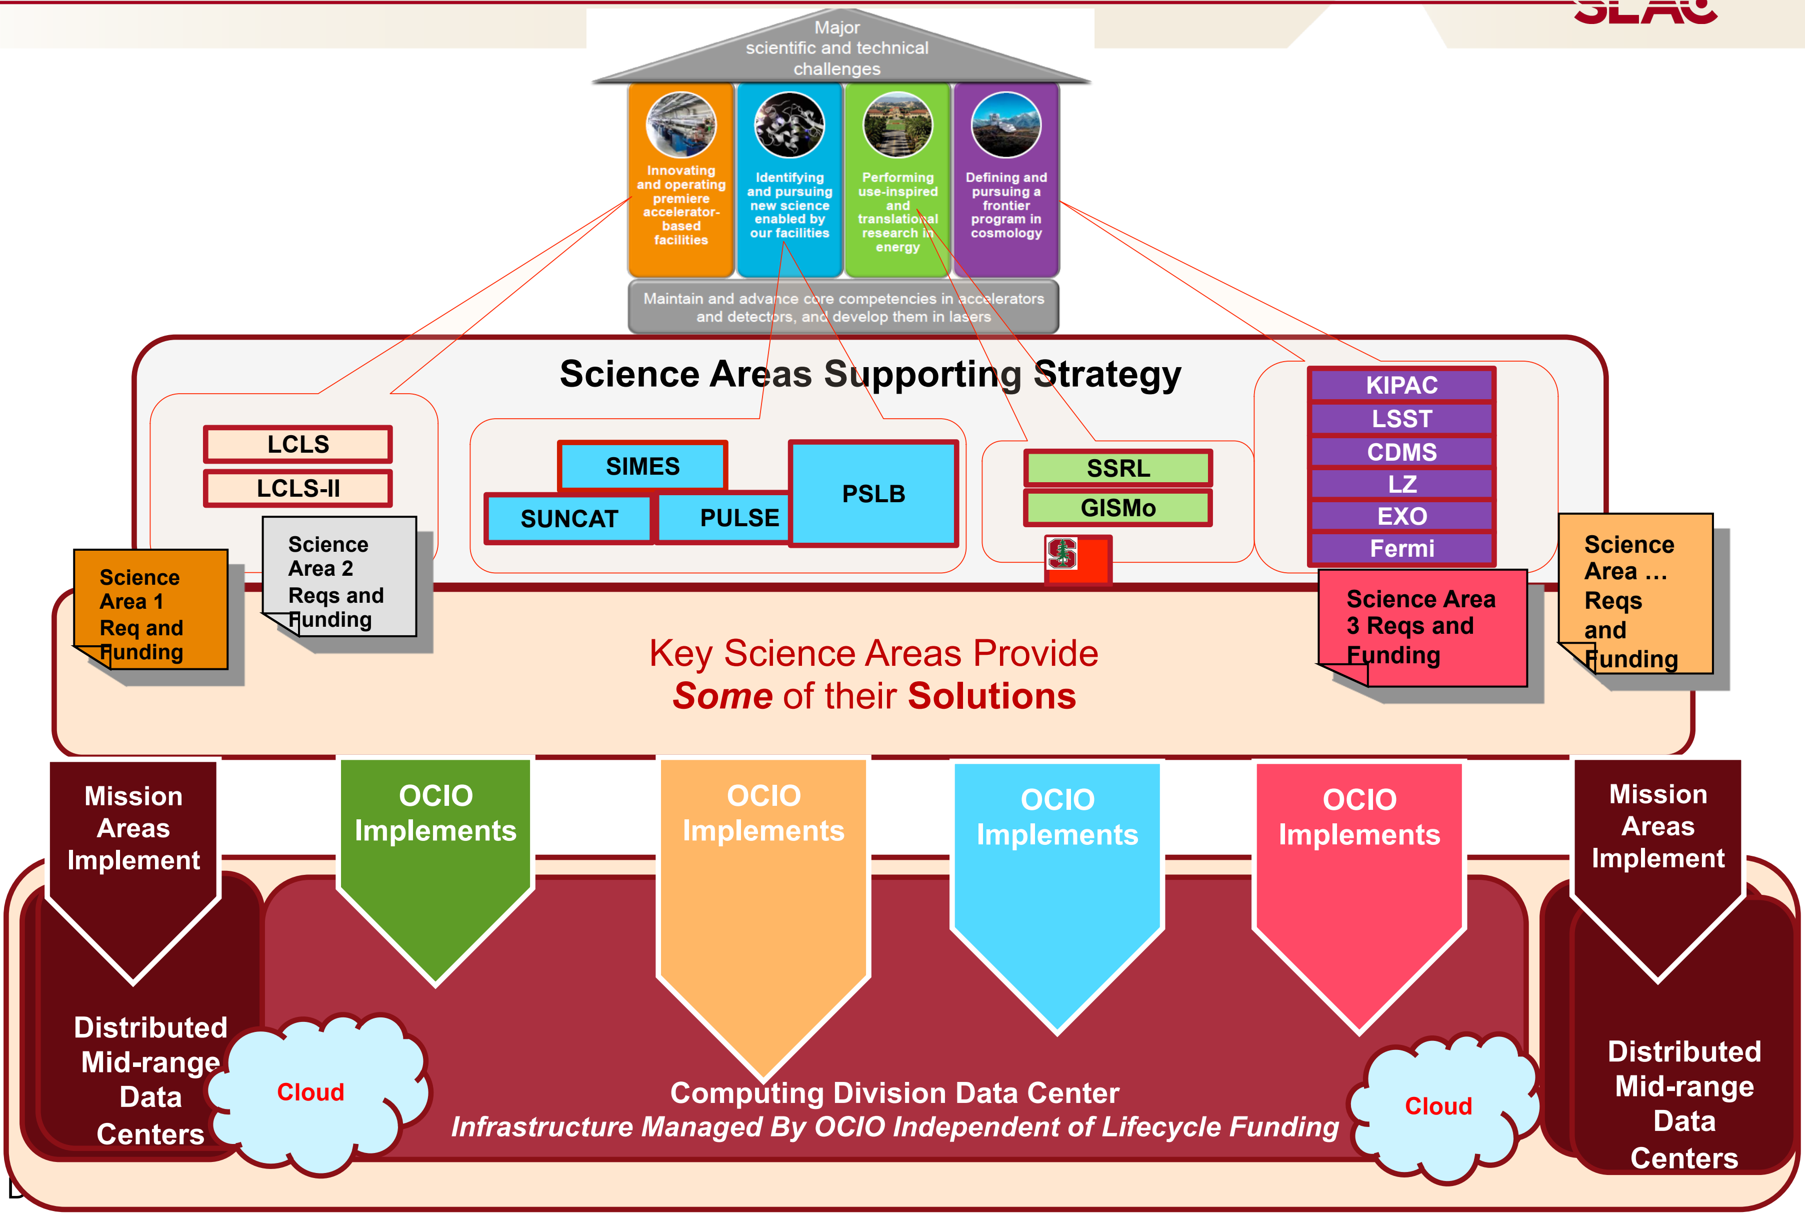

Proposed Paradigm to Gathering Integrated Requirements Across Lab as part of Computing Strategy

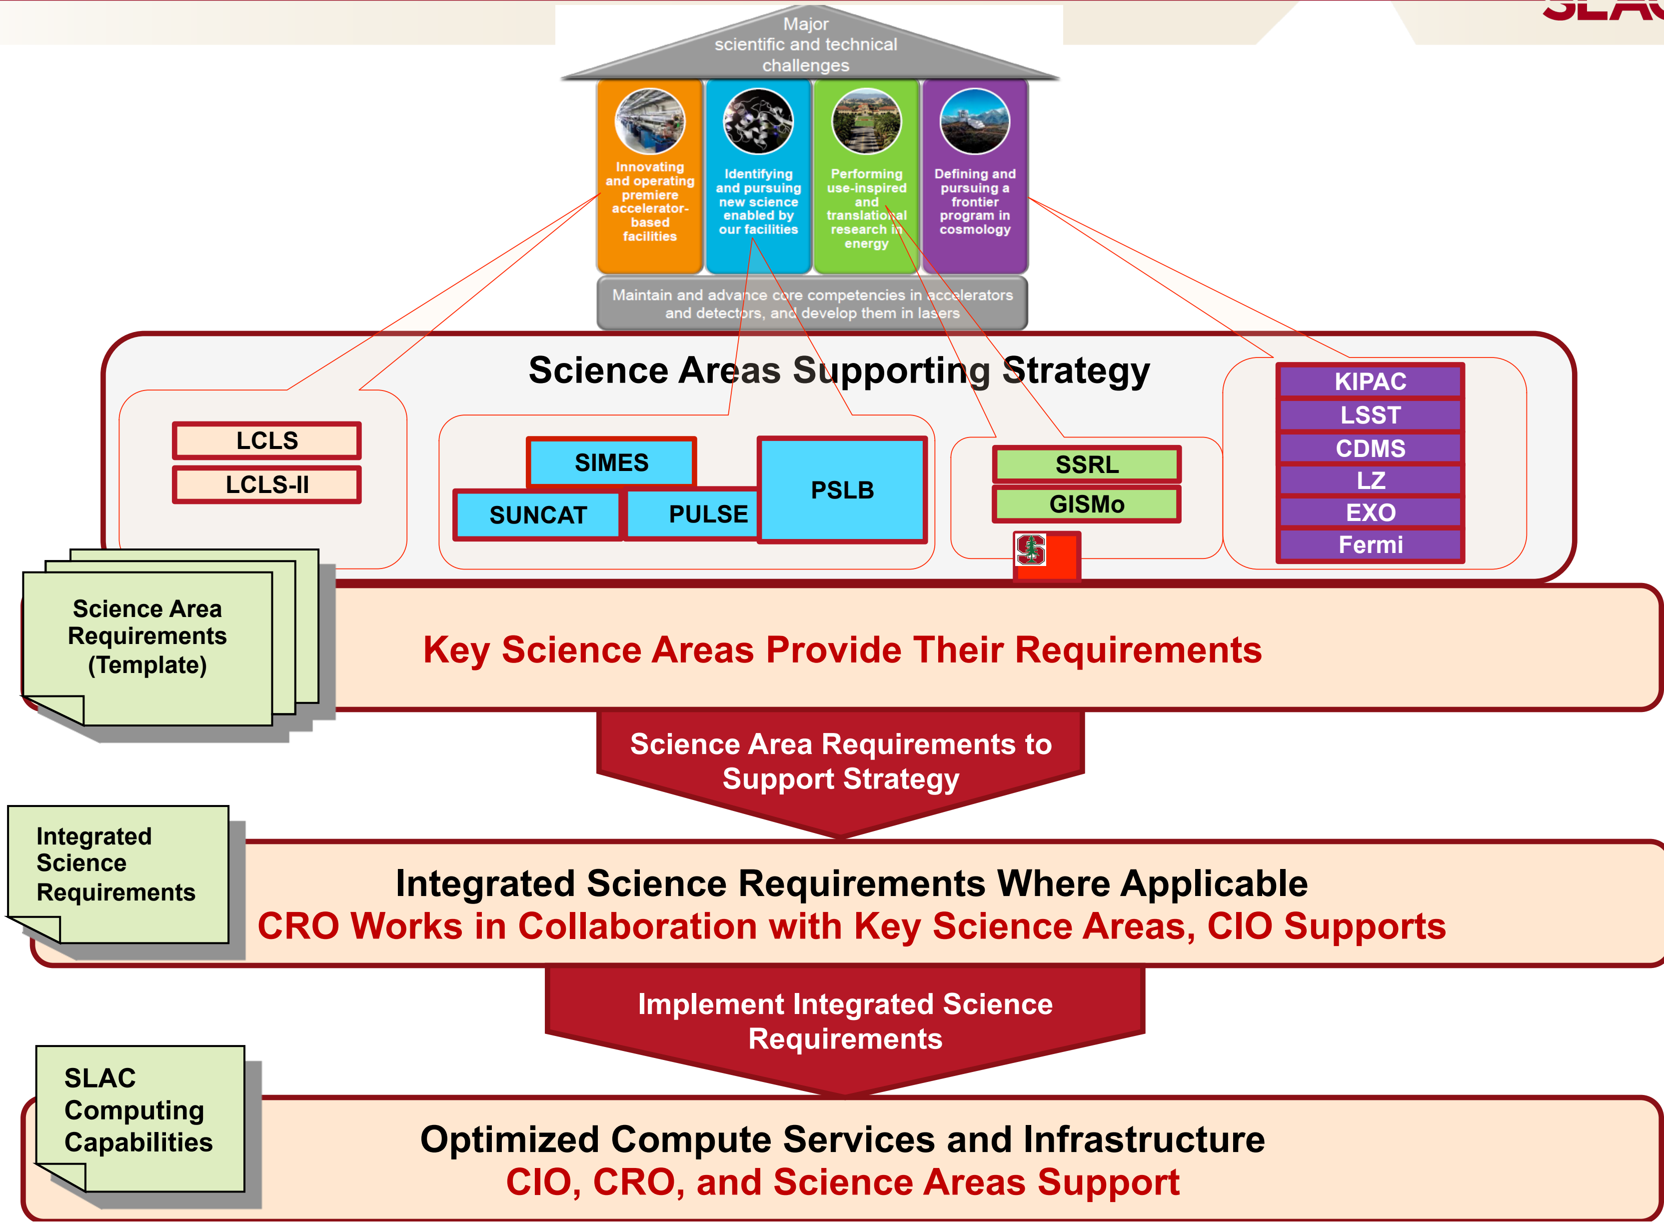

Scientific Computing “Continuum”

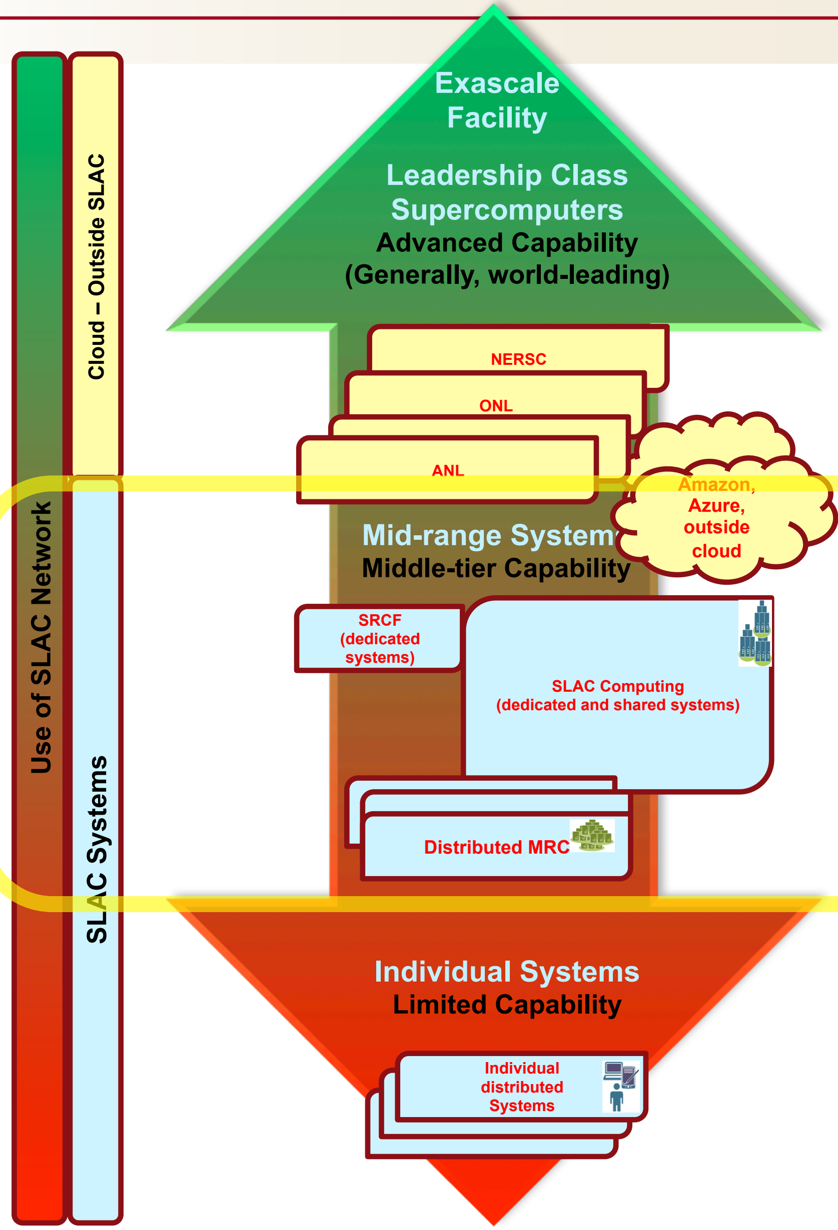

SLAC Mid-range Computing: Greatest Area of Opportunity

Varying resources (systems and people) scattered throughout Lab with major challenges associated with underutilization, duplication, lack of funding, and limited lifecycle management

IS3C Program Review

- Assembled a group of IT leaders from research and industry to review our scope and approach
- Representatives from Fermilab, NERSC, Brookhaven, Argonne, LLNL, PNNL, NASA, JPL
- We are the first computing organization to adopt this holistic scientific computing approach
- Reviewers provided valuable, constructive feedback

IS3C Summary

- Develop a complete view of end-to-end Scientific Computing requirements
- Consider all supporting resources: services, staffing, infrastructure and tools
- Identify any resource gaps
- Track resource metrics, especially data on infrastructure and facilities in support of Scientific Computing
- Identify commonality across requirements
- Develop sustainable business models for baseline computing capabilities
- We need your input (requirements capture, feedback)

Questions?

SLAC<->NERSC Partnership

Scientific Computing Services

Yemi Adesanya, March 2, 2017

Holistic view of Scientific Computing

- SLAC Scientific Computing may involve:
  - Laptops and PCs
  - Local mid-range clusters
  - Cloud
  - Supercomputers at leading facilities
  - Software applications and tools
  - Experts that can assist in developing and debugging solutions

- OCIO has a role to play as a facilitator to enable SLAC scientific computing
  - Reach out to the user community
  - Gather requirements and feedback
  - Deliver a cohesive user experience
  - Ensure SLAC users can leverage external resources effectively

SLAC-NERSC Partnership

- SLAC and NERSC support the same science! We are part of the same mission

- We must help SLAC users make effective use of NERSC:
  - Provide communication on NERSC resources and roadmap
  - Advise users on how to leverage NERSC effectively
  - Optimize SLAC network infrastructure and security to support SLAC-NERSC distributed computing and data management
  - Align SLAC mid-range compute to NERSC services

- Let’s identify any areas that have potential for collaboration

Storage & Data Management

Scientific Computing Services

Lance Nakata, March 2, 2017

U.S. DEPARTMENT OF ENERGY
Office of Science

SLAC NATIONAL ACCELERATOR LABORATORY

Storage Updates

- 15 T10000D 8TB tape drives now in HPSS production
  - Tape drives attached to 5 fast, SSD-based servers
  - astore/mstore data and many 5TB tapes migrated to 8TB tapes
  - All HPSS data now written at 8TB capacity, not 1TB or 5TB
- 1 SSD-based server for AuriStor (AFS) service in production
  - 25TB of usable space; better performance; more on the way
- GPFS 3.5 to Spectrum Scale 4.1 upgrades have begun
  - GPFS 3.5 is end-of-support-life on 4/30/2017
  - bullet cluster already upgraded to 4.1; file servers next
  - Two-step upgrade process from 3.5 to 4.1 then later to 4.2 to reduce/eliminate scheduled downtime

Storage Updates (2)

- Storage as a Service (StaaS) upgrades
  - SSDs for faster metadata operations
  - Some possible SSD space for small data needs
  - Planning enhancements to Clustered NFS service
- tsm1 tape backup server upgrade
  - SSDs for database and storage cache
  - Move from 1TB to 5TB tape drives
- T10000E tape drive cancelled
  - IBM TS-series and LTO are future tape drive candidates
  - Still using T10000D till higher density drive available

End-Of-Life Storage Hardware

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

- Solaris 10 support will end 1/31/2018. Hardware phaseout will continue through 2017.

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

Questions?

Unix Platform

Scientific Computing Services

Andrew May & Christa Doane, March 2, 2017

U.S. DEPARTMENT OF ENERGY
Office of Science

SLAC NATIONAL ACCELERATOR LABORATORY

Unix Platform Update

- Red Hat Enterprise Linux and CentOS
- Chef configuration management
- FastX
- Monitoring RHEL7

Red Hat Enterprise Linux and CentOS

- How to get started with CentOS 7:

https://confluence/display/SCSPub/CentOS+7+and+Chef

- Unless RHEL 7 support is requirement for your application (typically server), CentOS 7 is preferred and recommended instead
- RHEL5 End of Life, March 2017 (RHEL 6 EOL 2020)
- CentOS 7 desktop, ITDS portfolio of supported apps
  - standard portfolio of productivity apps like we have for Windows: mail client, web browser, ssh, office suite, etc.
  - CentOS 7 Desktops – for personal productivity, not servers

Chef configuration management

- 60 nodes using Chef management
- Chef workflow via Automate now works (Feb 2017)
- Expand use in our group now that we have a workflow framework
- Return to existing cookbooks to provide needed functionality
- Still lacking NFS/GPFS automounter maps, but actively working on it

- Chef allows users to configure their own services
  - LSST is starting to write their own Chef Cookbook
- This is one important reason we choose to move away from Taylor (monolithic in terms of adding/modifying configurations)

- Chef Premium features
- Automate: prescriptive workflow (continuous deployment pipeline, git version control, automatically test all changes on full VMs)
- Visibility: - dashboards, reports
- Compliance: - write compliance rules, view reports

FastX

- Display remote Linux applications (X Clients) on your desktop or laptop
- You can run FastX in any standard web browser
- There is also a standalone client application
- The FastX backend is hosted on a cluster of VMs in VMware
  - As demand necessitates, more VMs will be added
- Service has been quite stable the last 3 months
- Documentation to get started: https://confluence.slac.stanford.edu/pages/viewpage.action?pageId=205985167

Future System and App Monitoring

- Move to RHEL 7 – opportunity to implement new monitoring architecture
  - Goals
    - Continue utilizing Open Source products
    - Leverage modern and flexible approach
    - Possibilities for self service in graphing and alerting
  - Looking for input from science groups about important measurements or capabilities desired
    - Contact Christa Doane – cdoane@slac.stanford.edu

What would new architecture look like?

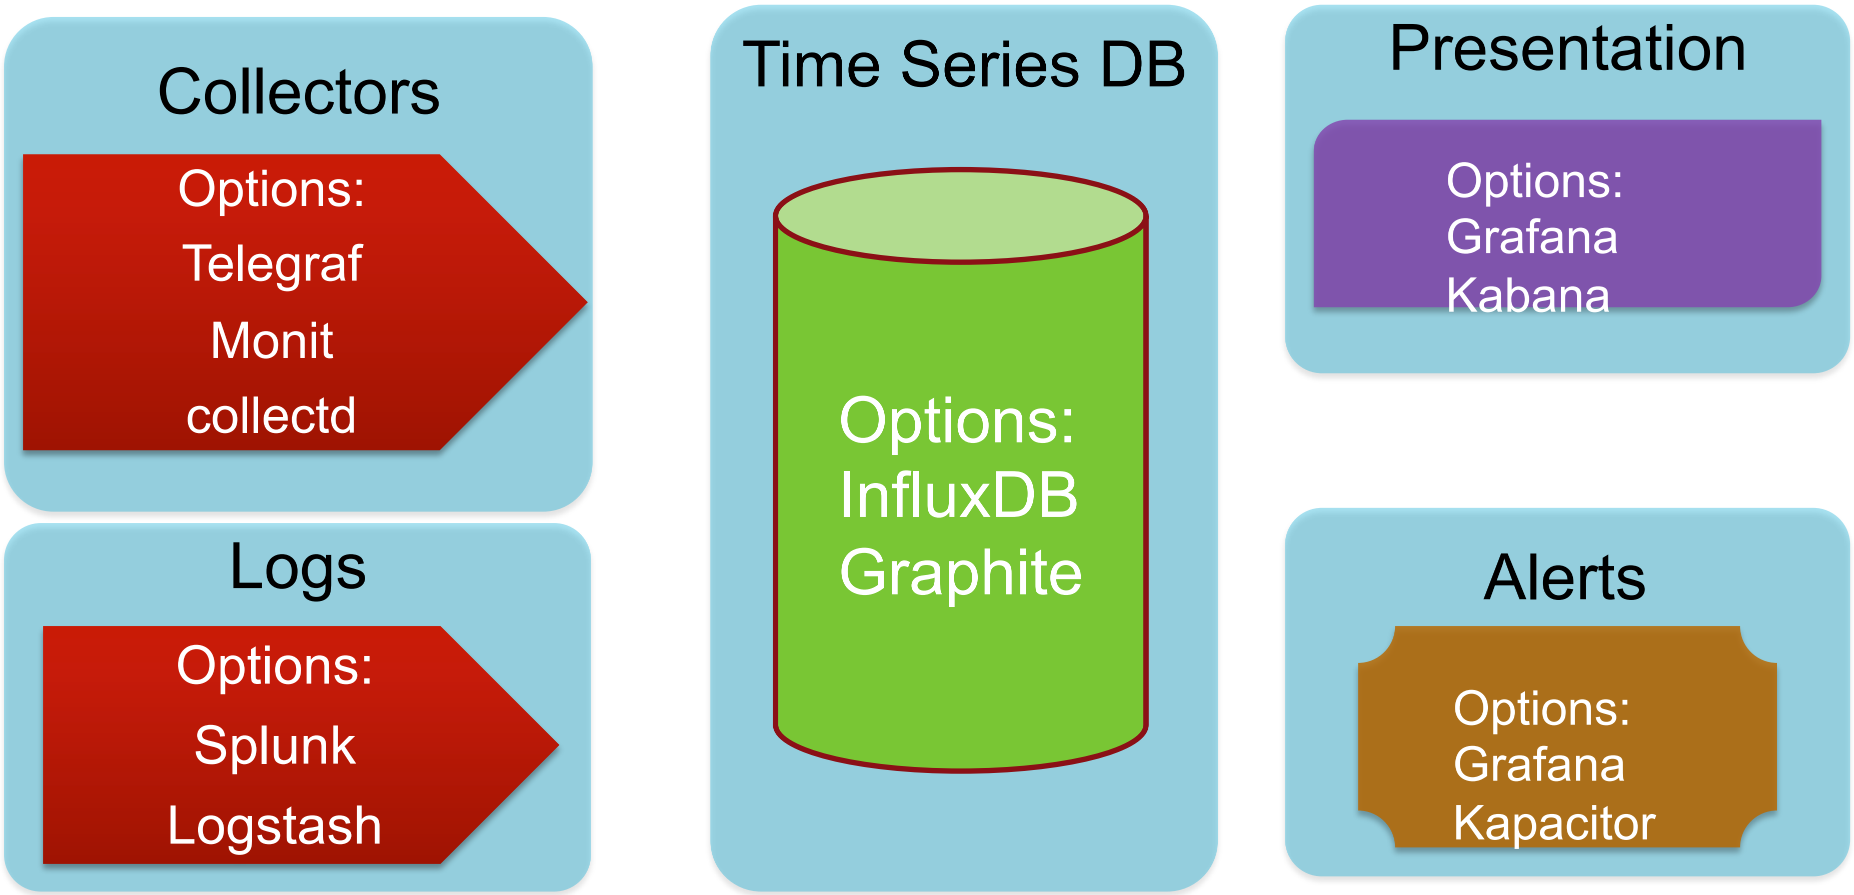

Questions?

Container Technology

Scientific Computing Services

Wei Yang, March 2, 2017

U.S. DEPARTMENT OF ENERGY
Office of Science

SLAC NATIONAL ACCELERATOR LABORATORY

What we will and will not do

Will do:

- Provide infrastructure to run containers
- For short development - OpenStack
- For larges scale operation, in batch or via OpenStack
- May provide several flavor of containers technologies, for now Docker

Will not:

- Develop customized container images for users

(Except perhaps generic ones for batch jobs, or specialized such as Jupyter via OpenStack)

We need:

- Initial use cases to drive this forward
- Work with SCS to define typical ways of using containers at SLAC
- Work with SCS to sort out issues

Use case 1: Geant 4

- Geant4 on Microsoft AZURE & SLAC
- Andrea Dotti as PI for Microsoft grant to run Geant4 in AZURE
- Will use Docker at AZURE
  - From scratch, simple use case first - use case that has a good chance to adapt to the environment.
- Also want to run at SLAC to verify the portability of the Docker image
  - Next version of LSF to be deployed at SLAC will support containers
  - Need input and output spaces. Temporarily from NFS
  - Docker image deployment with LSF will comply with SLAC security requirement

Use case 2: ATLAS experiment

- ATLAS is (finally) interested in containers - workshop at CERN on March 8
- Not sure what will happen. But if ATLAS can be a driving force if it can provide something for us to try

My Guess:

- CentOS 7 based
- Pre-packed software or depend on external CVMFS ?
- Require outbound TCP
- One Container image or several, can't be many.
  - How image update fits in our security requirement
- Each site has its own mechanism to access data.
  - Data includes "experiment data" in files, via http, via remote xroot access, and metadata
  - Will there be a site specific software hook?
- Will it come through the Grid CE ?

SLAC

Questions?