LCLS Network & Support Planning

Terri Lahey

Terri Lahey
lahey@slac.stanford.edu

Outline

- Goal: build production hosts, workstations, & networks
- Engineering Teams
- Apply experience and new architectures
- Integrated Security at SLAC
- Servers & desktops
- Network Plans
- Ethernet Architecture
- What's Next?

Terri Lahey
lahey@slac.stanford.edu

Engineering Teams

- Scientific Computing & Computing Services (SCCS) network and security:

  Gary Buhrmaster et. al., Antonio Ceseracciu,

  Charles Granieri, Fred Hooker
- LCLS:

  Mark Crane, Mike DiSalvo, Doug Murray
- Controls & Power Engineering (CPE):

  Ken Brobeck, Jim Knopf,

  Terri Lahey, Jingchen Zhou

Terri Lahey
lahey@slac.stanford.edu

Apply Experience from PEP and Implement New Architectures

- Protect accelerator components and access to the control system
  - Control number of connections
  - Control who connects
- Meet Users needs
  - Physicists, operators, engineers need access to control system and components so they can do their job
- Implement Security for the networks and hosts on the network

Terri Lahey
lahey@slac.stanford.edu

- Commission LCLS Injector from MCC control room
- Physicists, Engineers & Operators will use:
  - EPICS
  - Matlab
  - existing HLAs (SLC)

Terri Lahey
lahey@slac.stanford.edu

Use SCCS services where possible

- Security:
  - Work with SCCS security team to help us run 24x7
  - SCCS security coordinates SLAC-wide security
    - identify model and DOE/Office of Science requirements
    - Interfaces with DOE/Office of Science
    - Scan networks in a scheduled manner (production very controlled)
  - Participate in Computing Security Committee
- Network Design and Physical Layer
  - SLAC standards to achieve more reliable networks
  - Central Management with strong liaison to Controls
  - Current Equipment/Design Knowledge
- SCCS manages Oracle, web servers. Servers reside at MCC
- Use AFS for CVS repository, development, & main web server (mirror to MCC).
- Use SCCS central tools when possible: console log management, authentication

Terri Lahey
lahey@slac.stanford.edu

Production Servers & Workstations

- Manage production servers to run standalone
- Use SCCS-supported versions of operating systems, packages & applications where possible
- Patch operating systems and update to new versions
- Automate maintenance of production hosts
- Reduce maintenance load and improve security by using taylor where possible
- Centralized Log server & security monitoring
- Use existing accelerator production servers where possible (e.g. NFS,elog, ARTEMIS bug tracking, ORACLE, DNS, IP Services)

Terri Lahey
lahey@slac.stanford.edu

Networks

- SCCS Networking configures the network switches and routers & manages the physical layer.
- Controls Software coordinates control system and user needs, and works closely with SCCS.
- Production accelerator network is controlled and protected.
  - Greater attention to security by both SCCS and Controls
  - Run accelerator disconnected from the rest of SLAC; For use if there is a security problem at SLAC.
- Isolation of Wireless network:
  - Wireless and Accelerator switches are never combined.
  - Wireless is visitornet that resides outside SLAC firewall.
  - Users tunnel into SLAC the same way they tunnel from internet: ssh, citrix, vpn

Terri Lahey
lahey@slac.stanford.edu

Networks (cont'd)

- CISCO switches and routers
- Patch network firmware and upgrade versions.
- Plan for and upgrade hardware components to avoid end-of-life
- Implement Redundancy in core switches and routers, for reliability. Use hot spares for device switches, but increased use of VLANs will likely require some configuration.
- SLAC-wide Network monitoring systems send alarms:
  - components go offline (e.g.. power outage or failure)
  - ports get disabled due to too many collisions
- Enhance network monitoring

Terri Lahey
lahey@slac.stanford.edu

Technology Choices

- Cisco switches - gigabit:
  - Device switches: 3750 (single and stacks)
  - Core: pair of 6509 (720Gbps bidirectional backplane) supporting uplinks and servers
  - MCC control room workstations, printers: 4506
  - Wireless: 3750 (10/100) public switch
- Linux & RTEMS
  - RHEL3 or RHEL4
  - DELL
- SUN Ray Thin Clients & some Linux workstations
- DIGI Terminal Servers

Terri Lahey
lahey@slac.stanford.edu

Network Architecture

- Production accelerator network is isolated:
  - Protect IOCs that often require insecure services like telnet/rsh or have less secure tcp/ip stacks
  - Control access to accelerator components so that systems do not get overloaded
  - Use private addresses
  - Multiple VLANs to separate traffic
  - Ports disabled by default
  - 1gigabit to the end devices. Currently 1gigabit uplinks to MCC
- DMZ is only access to private network (login servers, web servers, PV gateways).
- MCC and SLC-aware IOC uses PEP proxy server
  - have tested with PEP running
  - 9 SLC-aware IOCs for injector
  - more testing to confirm that PEP & LCLS will not impact each other.
- path to SCCS data silos & other required sevices

Terri Lahey
lahey@slac.stanford.edu

LCLS Accelerator Network

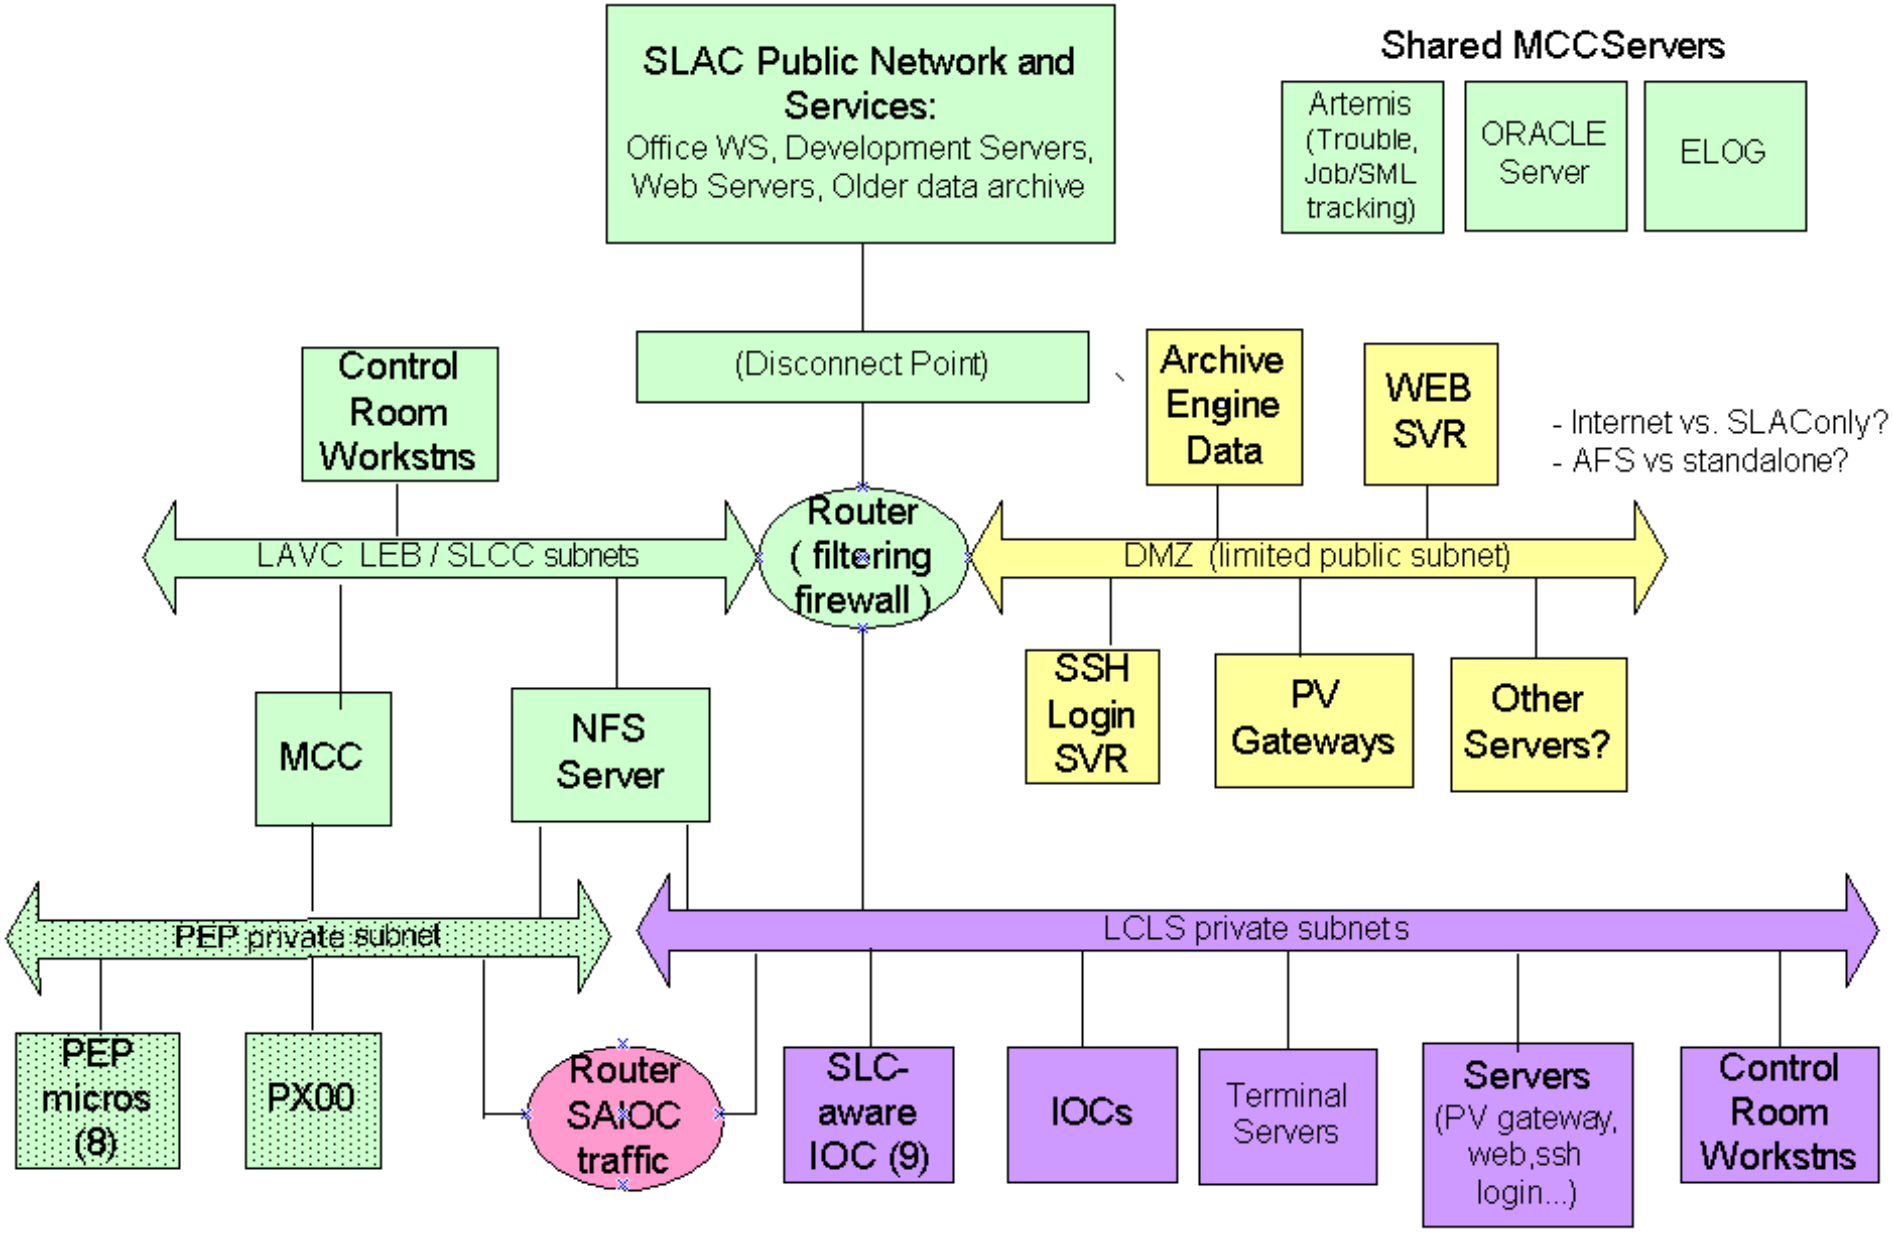

Current Work

- Building Production Infrastructure for Injector Commissioning Jan 2007
  - Installing network infrastructure in S20 & MCC
  - Additional tests of SLC-aware IOC and improving monitoring of traffic to avoid interference between PEP & LCLS programs
  - Review and implement network VLANs
  - Testing RHEL4 and working on production hosts
  - Ordered SUN Ray & will test during this PEP run
  - Integration with all LCLS subsystems

Terri Lahey
lahey@slac.stanford.edu

Conclusion

- Would like to hear your experiences:
  RHEL4, EPICS traffic, any isolated networks, archive data storage/management
- What worked well & what did not?

Terri Lahey
lahey@slac.stanford.edu

Thank you

Terri Lahey
lahey@slac.stanford.edu