Controls Computing and Networking Infrastructures for TestFac, FACET-II and LCLS/LCLS-II

Jingchen Zhou

Mission

SLAC has large, highly distributed and heterogeneous controls infrastructures with mixture of various legacy systems for multiple accelerator operations. The infrastructures have been evolving and growing with complexities and tight budget, and must be compliant with DOE Cyber Security. The system team:

- striving to tackle such a challenge and to provide reliable and secure controls networking and computing infrastructures highly critical to SLAC accelerator operations for
  - Test Facilities
  - FACET
  - LCLS and LCLS-II
- striving to perform systematic analysis based on all constrains and interfaces, and deliver sustainable solutions to meet growing demands
- committed to providing technical support to address issues and challenges constantly posed by OPS, engineers, and physicists
- leading efforts in networking and computing Infrastructure design for accelerator controls

System team members: Charles Granieri, Arjun Shetty, Ken Brobeck, Jingchen Zhou

Computing Infrastructure Support for Test Facilities

- Controls Computing Infrastructure for Test Facilities described in

  http://www.slac.stanford.edu/grp/cd/soft/unix/slaconly/testfac.html
- The infrastructure designed and implemented to support all Test Facilities programs in a centralized fashion and with the consideration of tight budget.

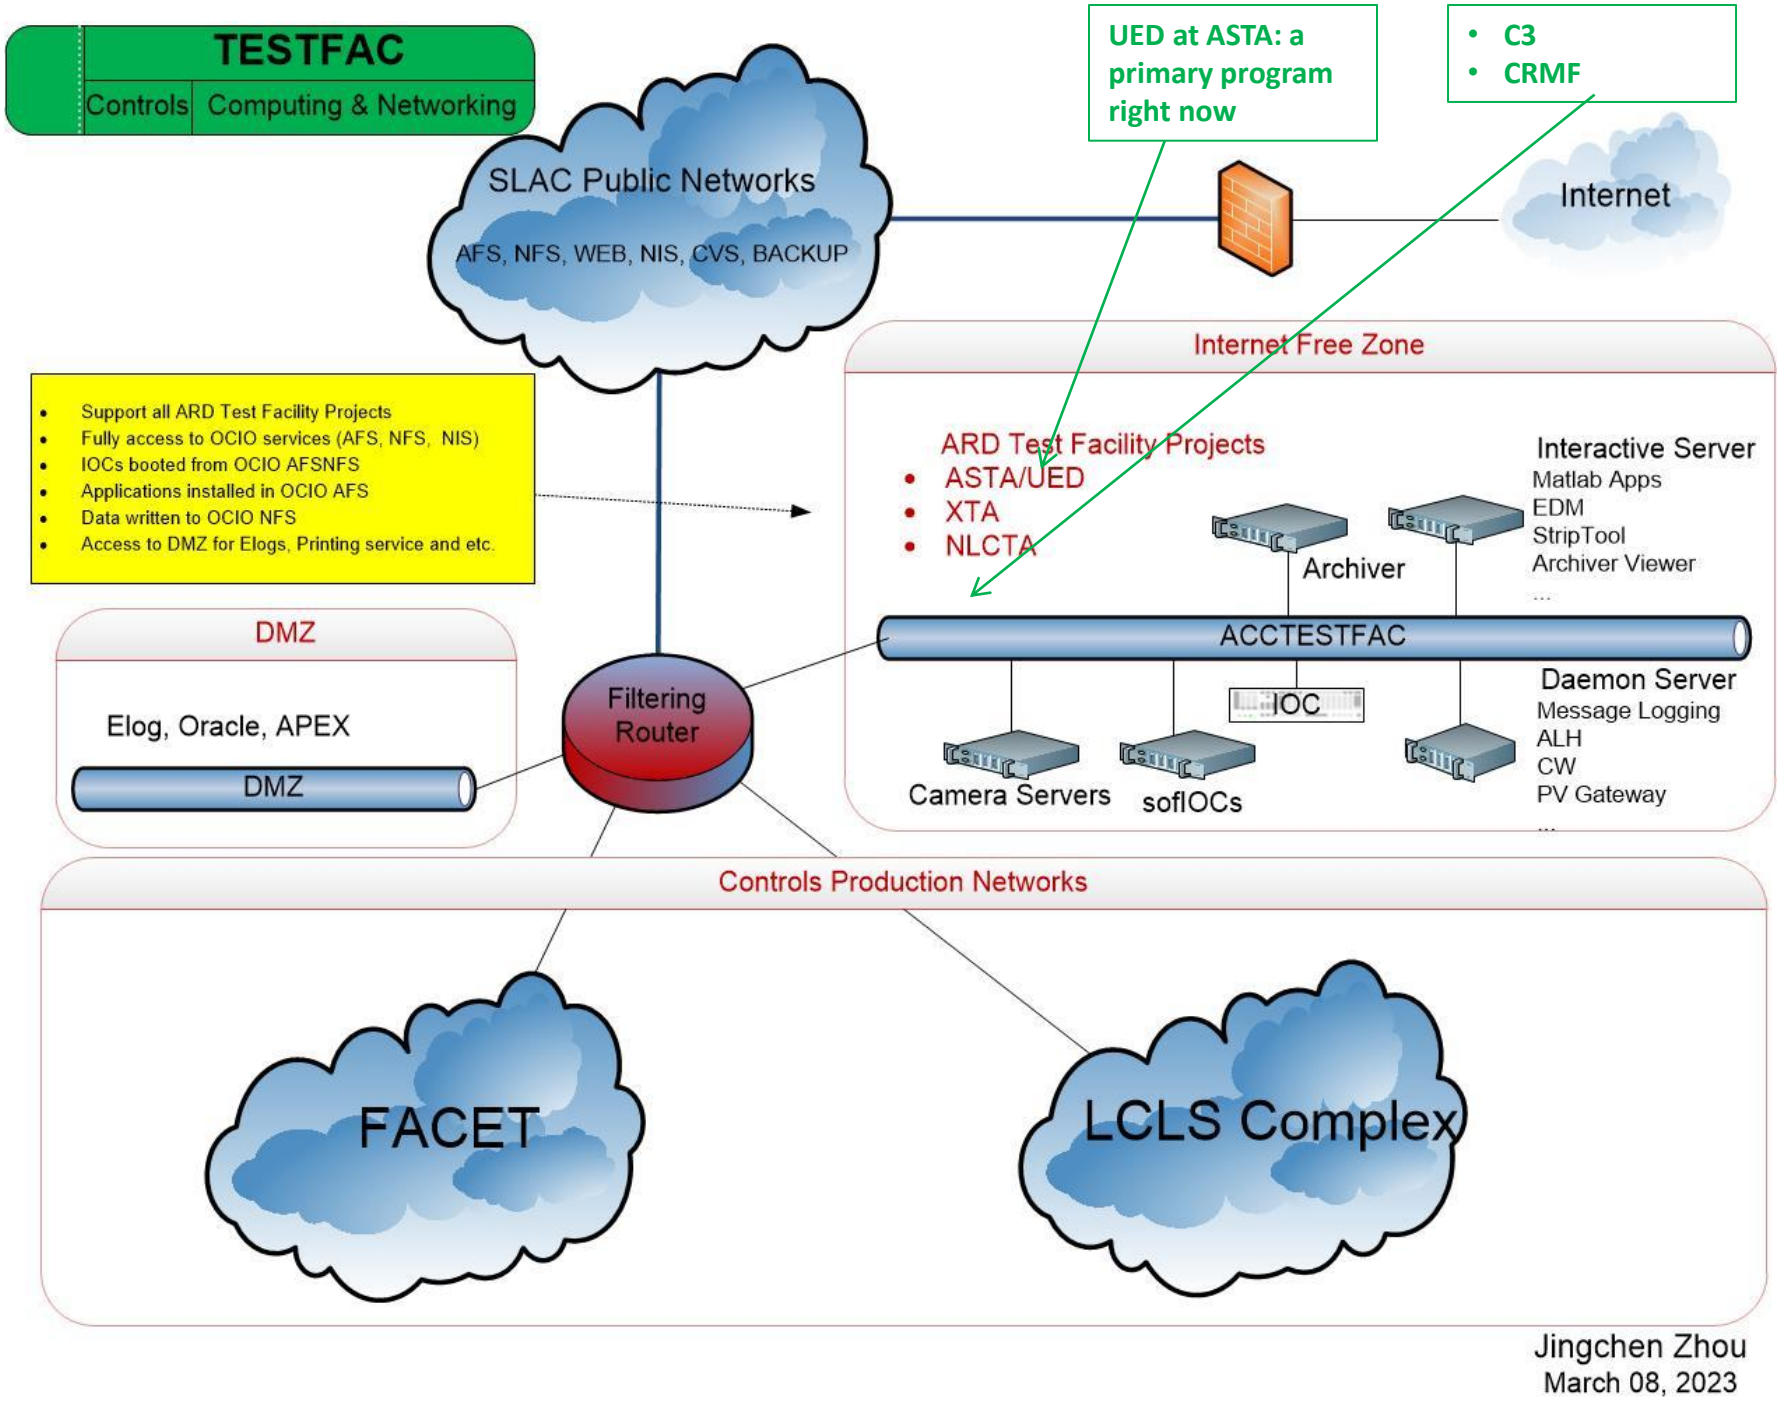

Jingchen Zhou
March 08, 2023

UED at ASTA

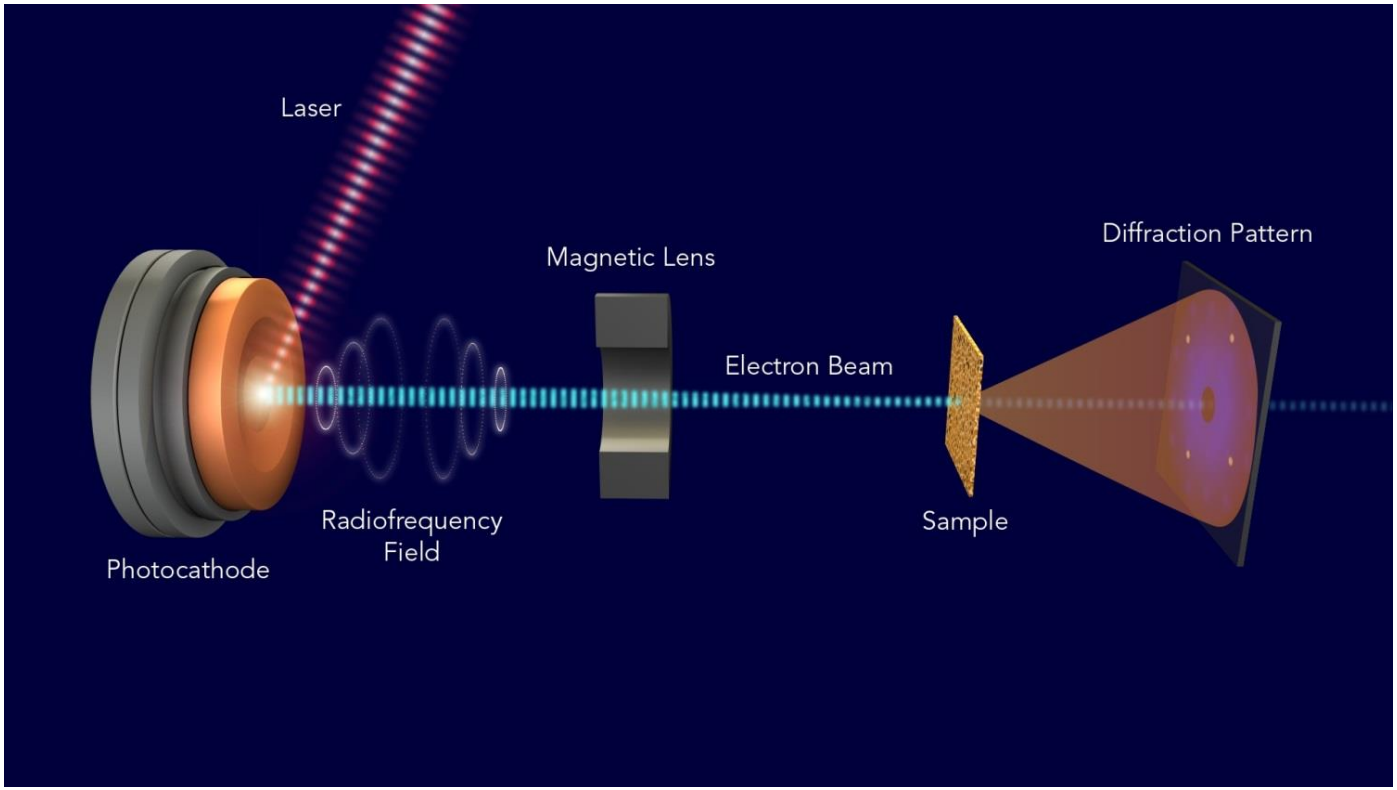

Image from Dr. X J Wang

- Ultrafast Electron Diffraction based on ASTA
  - Electrons scatter off the sample's atomic nuclei and electrons, creating diffractions.
  - Changes in the diffraction images over time are used to reconstruct the ultrafast processes, thus for the study of time-resolved, ultrafast atomic & molecular dynamics
- Complementary to the ultrafast studies with LCLS X-ray
  - Electrons interact differently with materials and "see" different things (i.e., reveal different properties) than Photons
  - Combined to draw a more complete picture of ultrafast processes within materials in complex systems

TestFac Production Environment

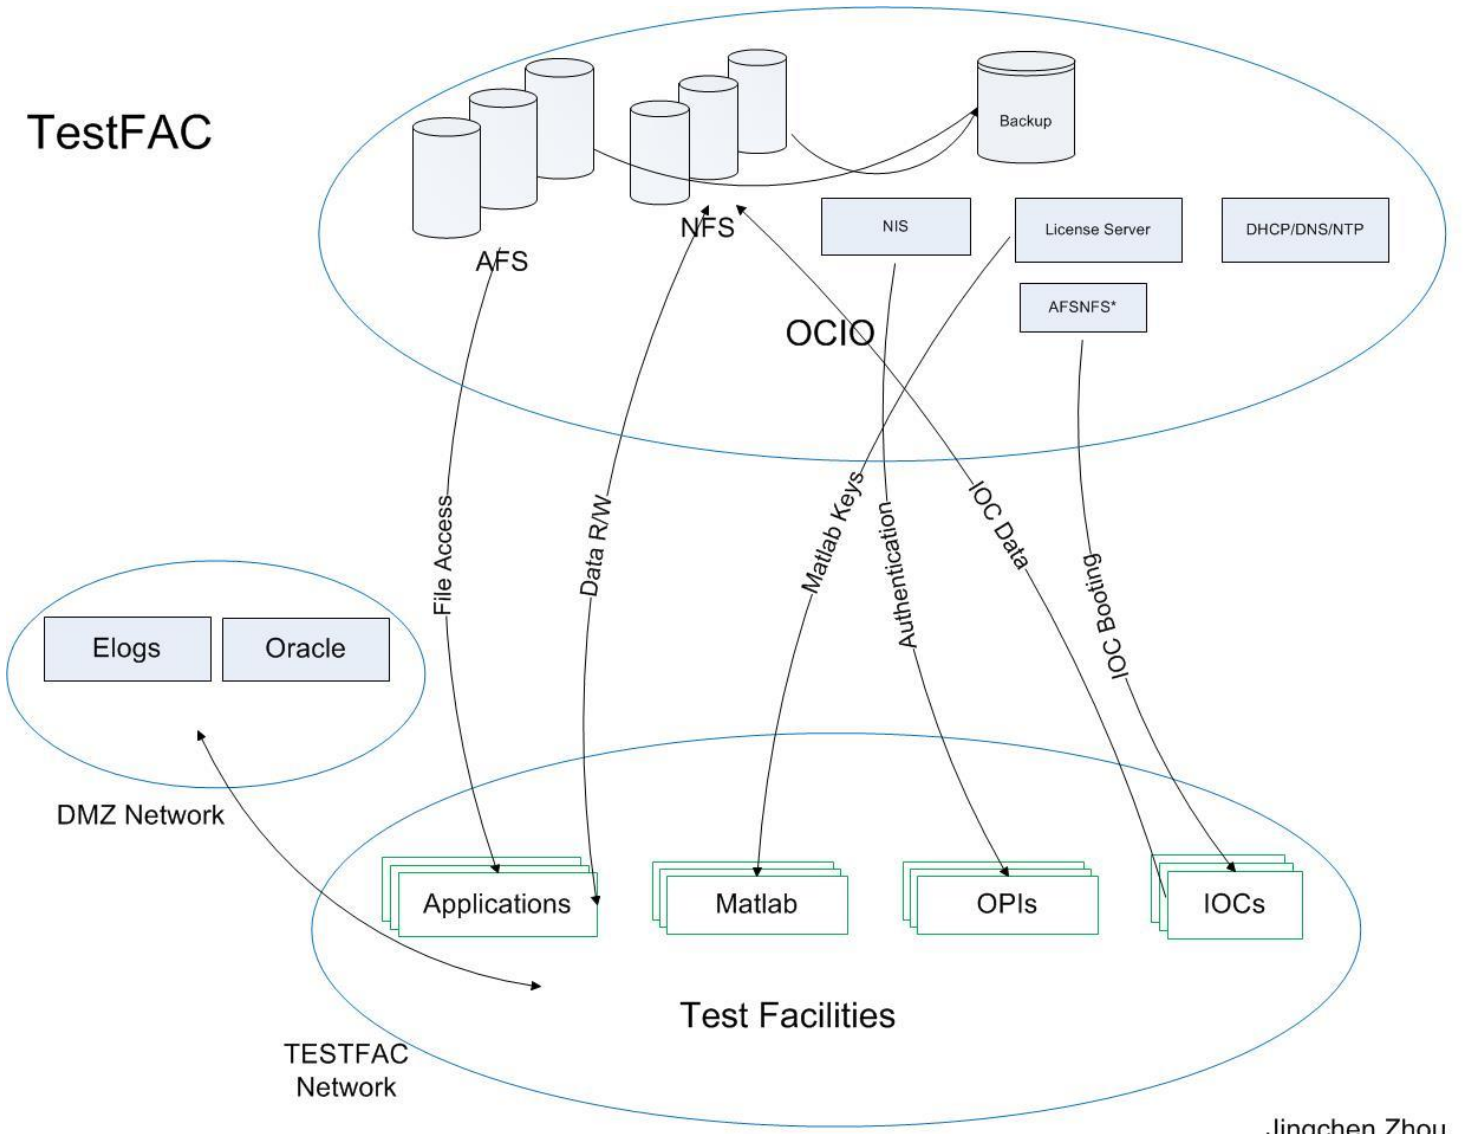

Jingchen Zhou
Friday, March 25, 2016

Summary

- This is a small but complete infrastructure with growing complexity for Test Facilities programs
  - C3

    https://slac.sharepoint.com/:p:/r/sites/controls/_layouts/15/Doc.aspx?sourcedoc=%7BF0886851-CE5F-4C6B-9FF8-FF7F3C5780CC%7D&file=C3%20Demo%20Controls.ppt&action=edit&mobileredirect=true
  - CRMF
- The infrastructure is critical to the success of Test Facilities programs
- The on-going effort is required for sustainability and being compliant with DOE Cyber Security

FACET

FACET Accelerator

FACET Experiments (S20)

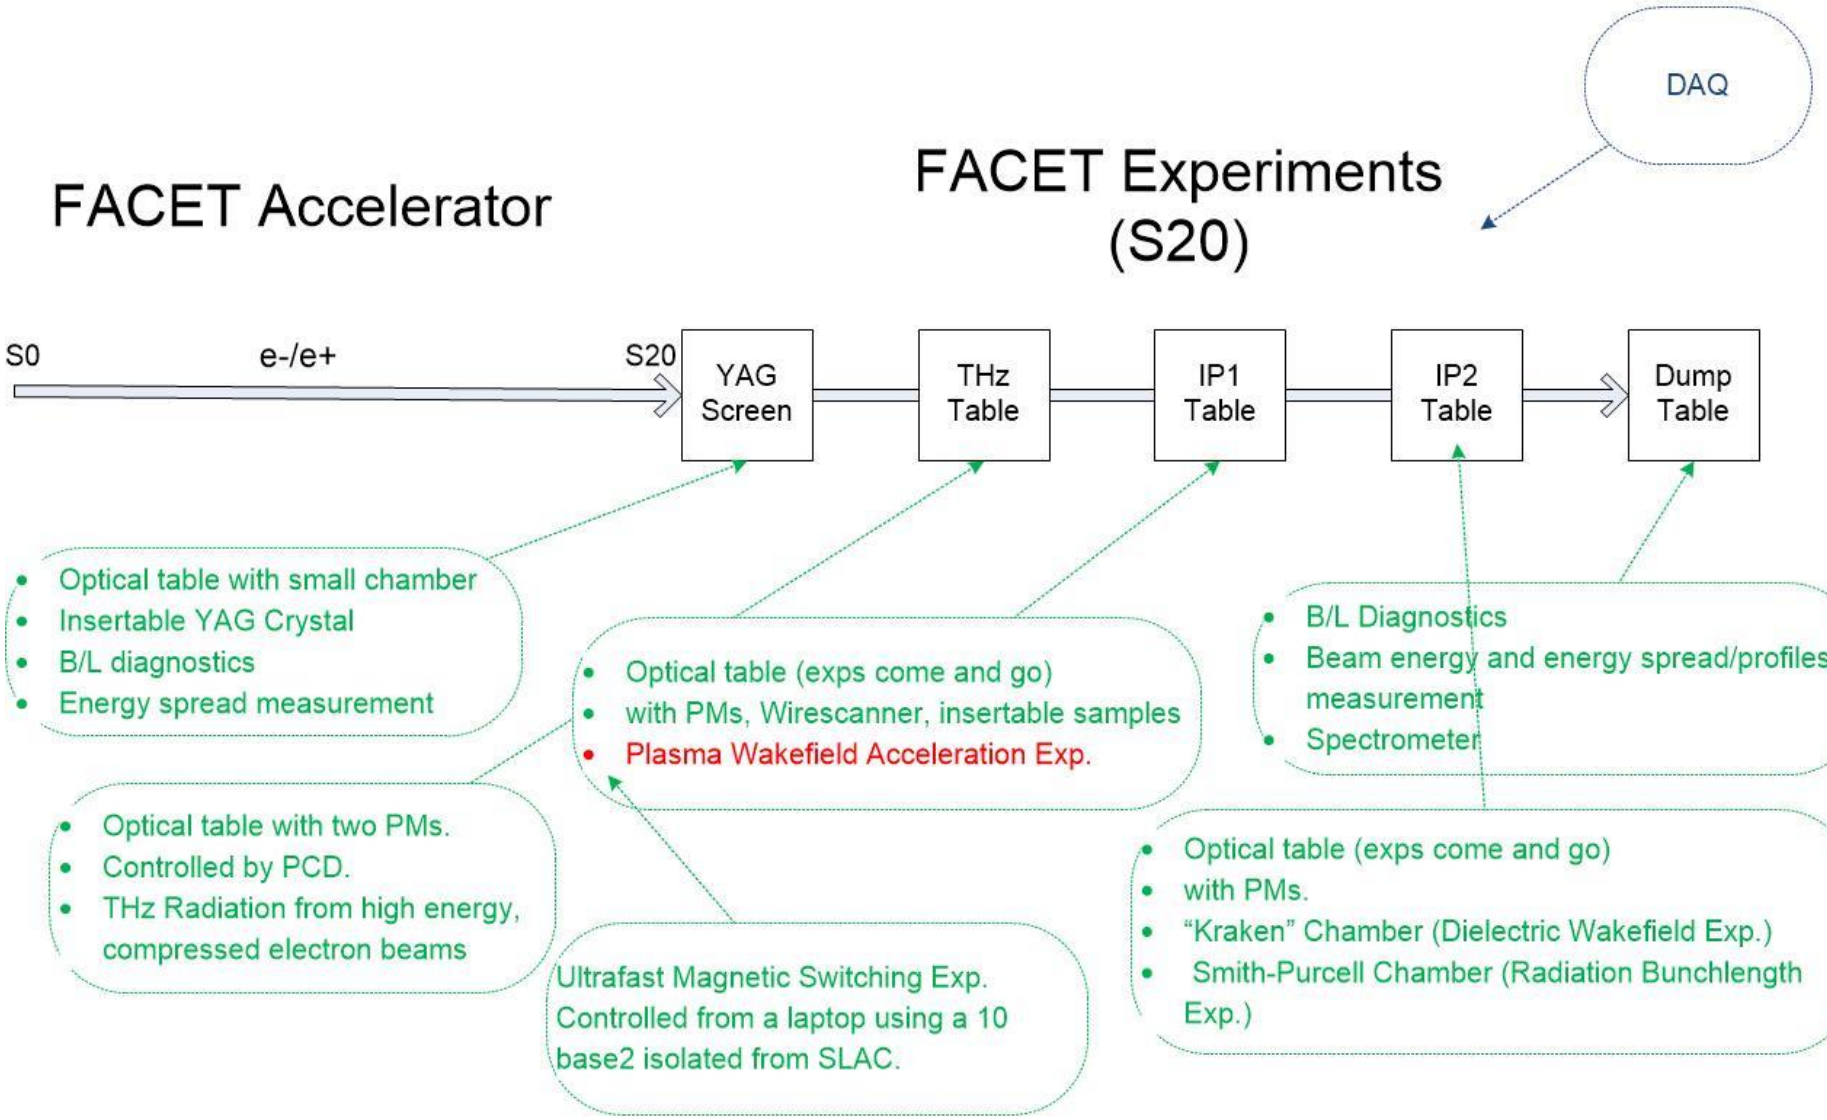

Jingchen Zhou

The Plasma Wakefield Accelerator

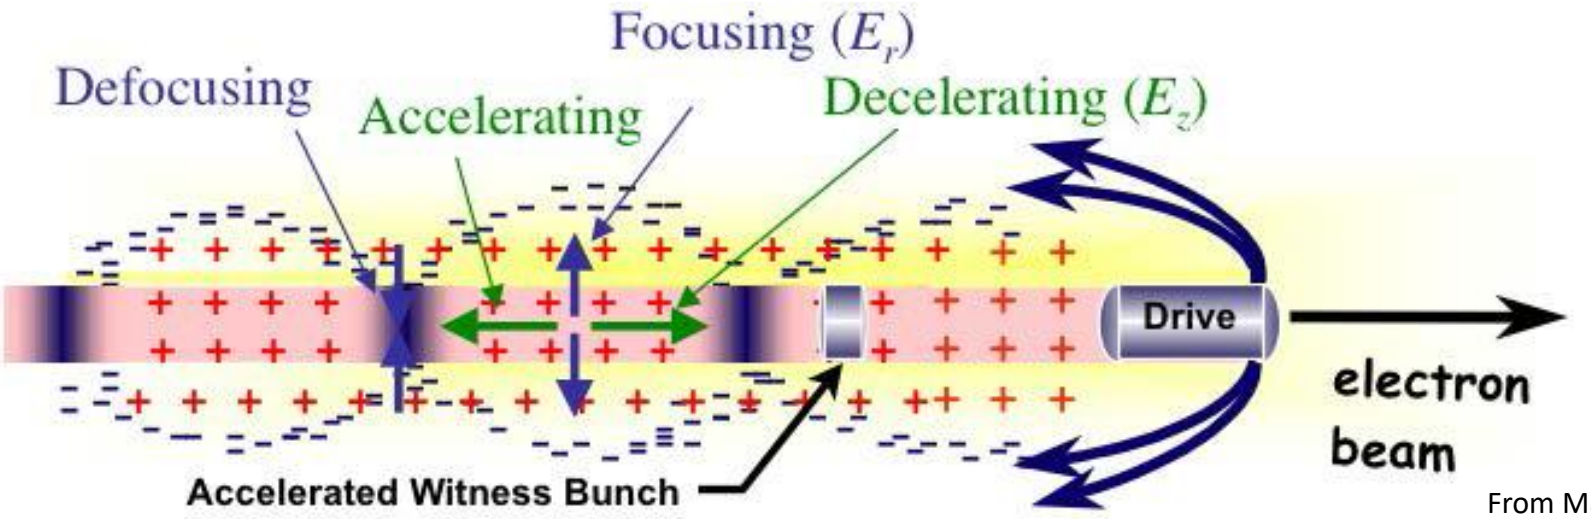

From Mark Hogan

Beam Driven

- Plasma Wakefield excited by relativistic e- bunch
- Tailing (witness) bunch accelerated by the Wakefield

Wakefield:

- None linear: Ez (accelerating, decelerating); Er (focusing, defocusing)
- high gradient (~ GeV/m)

FACET-II : optimization of e- acceleration (low emittance, narrow energy spread)
high gradient e+ acceleration

FACET Exp. DAQ

- Constraints: Controls Computing and Networking infrastructure designed for LCLS and FACET accelerator controls. Not adequate for experiment DAQs. NFS: serving accelerator controls apps, not a storage for incremental data at high rate. Controls Networking not adequate for high rate DAQ data. Security. No impact to LCLS operation.

- A patch but with some limitations:
DAQ data must be migrated to the public for analysis.
NO apps running on accelerator networks allowed to access DAQ data (both images and CA large waveforms) at high rate.
No dual-home allowed due to security constraint

- NAS recommendation: reliable, scalable, with performance.

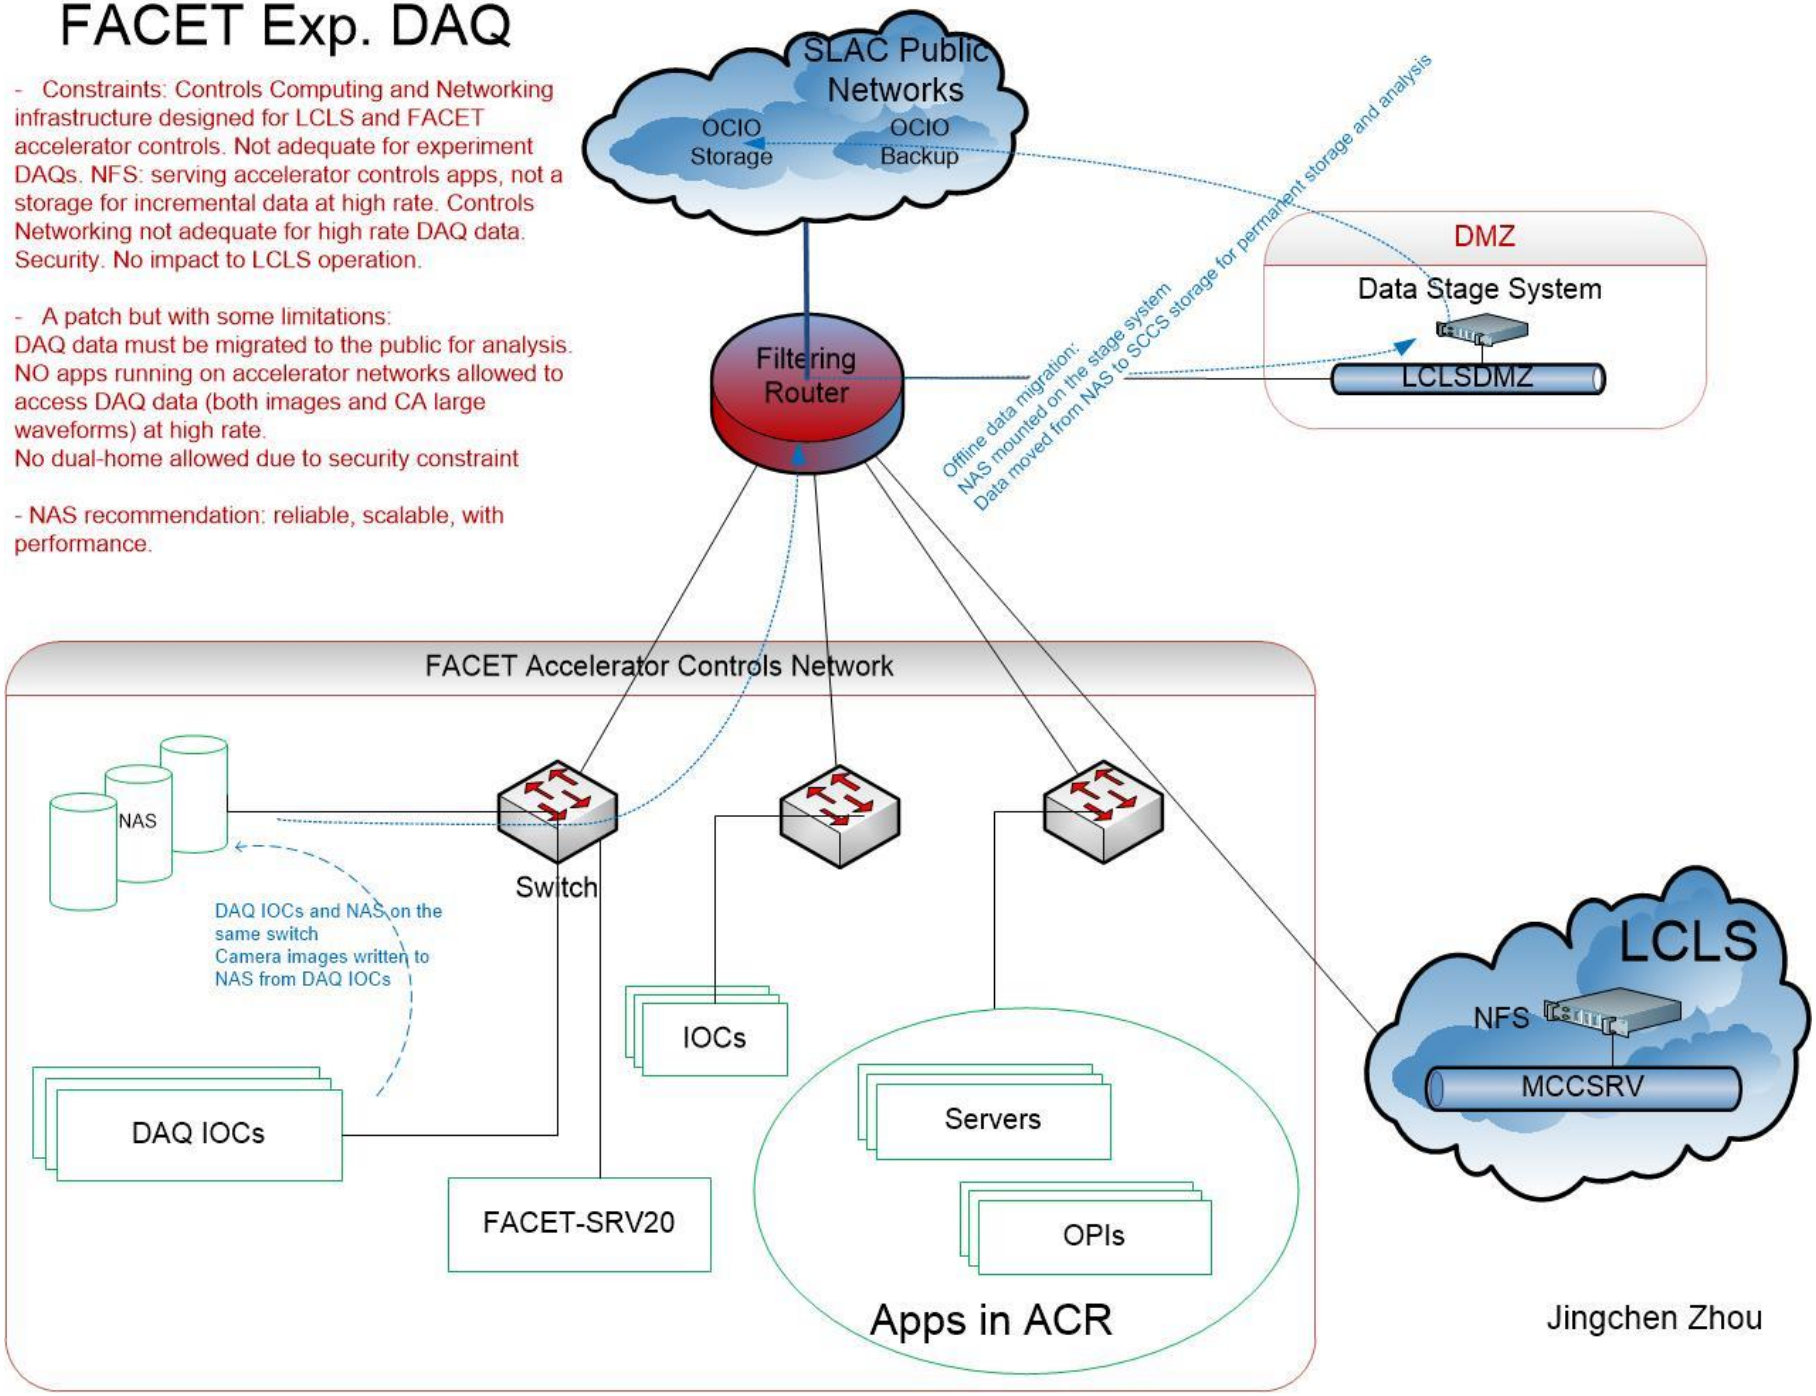

Jingchen Zhou

VMS

- FACET-II still depends on SLC controls, which is based VMS, a legacy system, also critical to A-line and NLCTA programs
- The VMS has been becoming unreliable and must be managed properly and proactively to keep it alive. Special thanks to Ken Brobeck and Ed Miller.
- If VMS were down, FACET-II would stop.
- Good news: VMS upgraded
  - Completed upgrading the OS to VSI 8.4.
  - Completed upgrading packages above the OS.
  - Completed testing SLC controls systems and EPICSv7 based data providers.

  The upgrades are critical to support FACET-II on VMS with sustainability and maintainability.

Summary

- The computing infrastructure for FACET-II
  - Highly distributed
  - Large with complexity
  - Standalone on the private network, but much parasited on LCLS
  - Mostly aging (i.e., requiring more maintenance)
- The beauty of this design enables SLAC to continue do the great science using legacy systems and with less money. Well aligned with SLAC mission!
- The on-going effort
  - Required for sustainability and being compliant with DOE Cyber Security
  - To provide a foundation for FACET-II to flourish

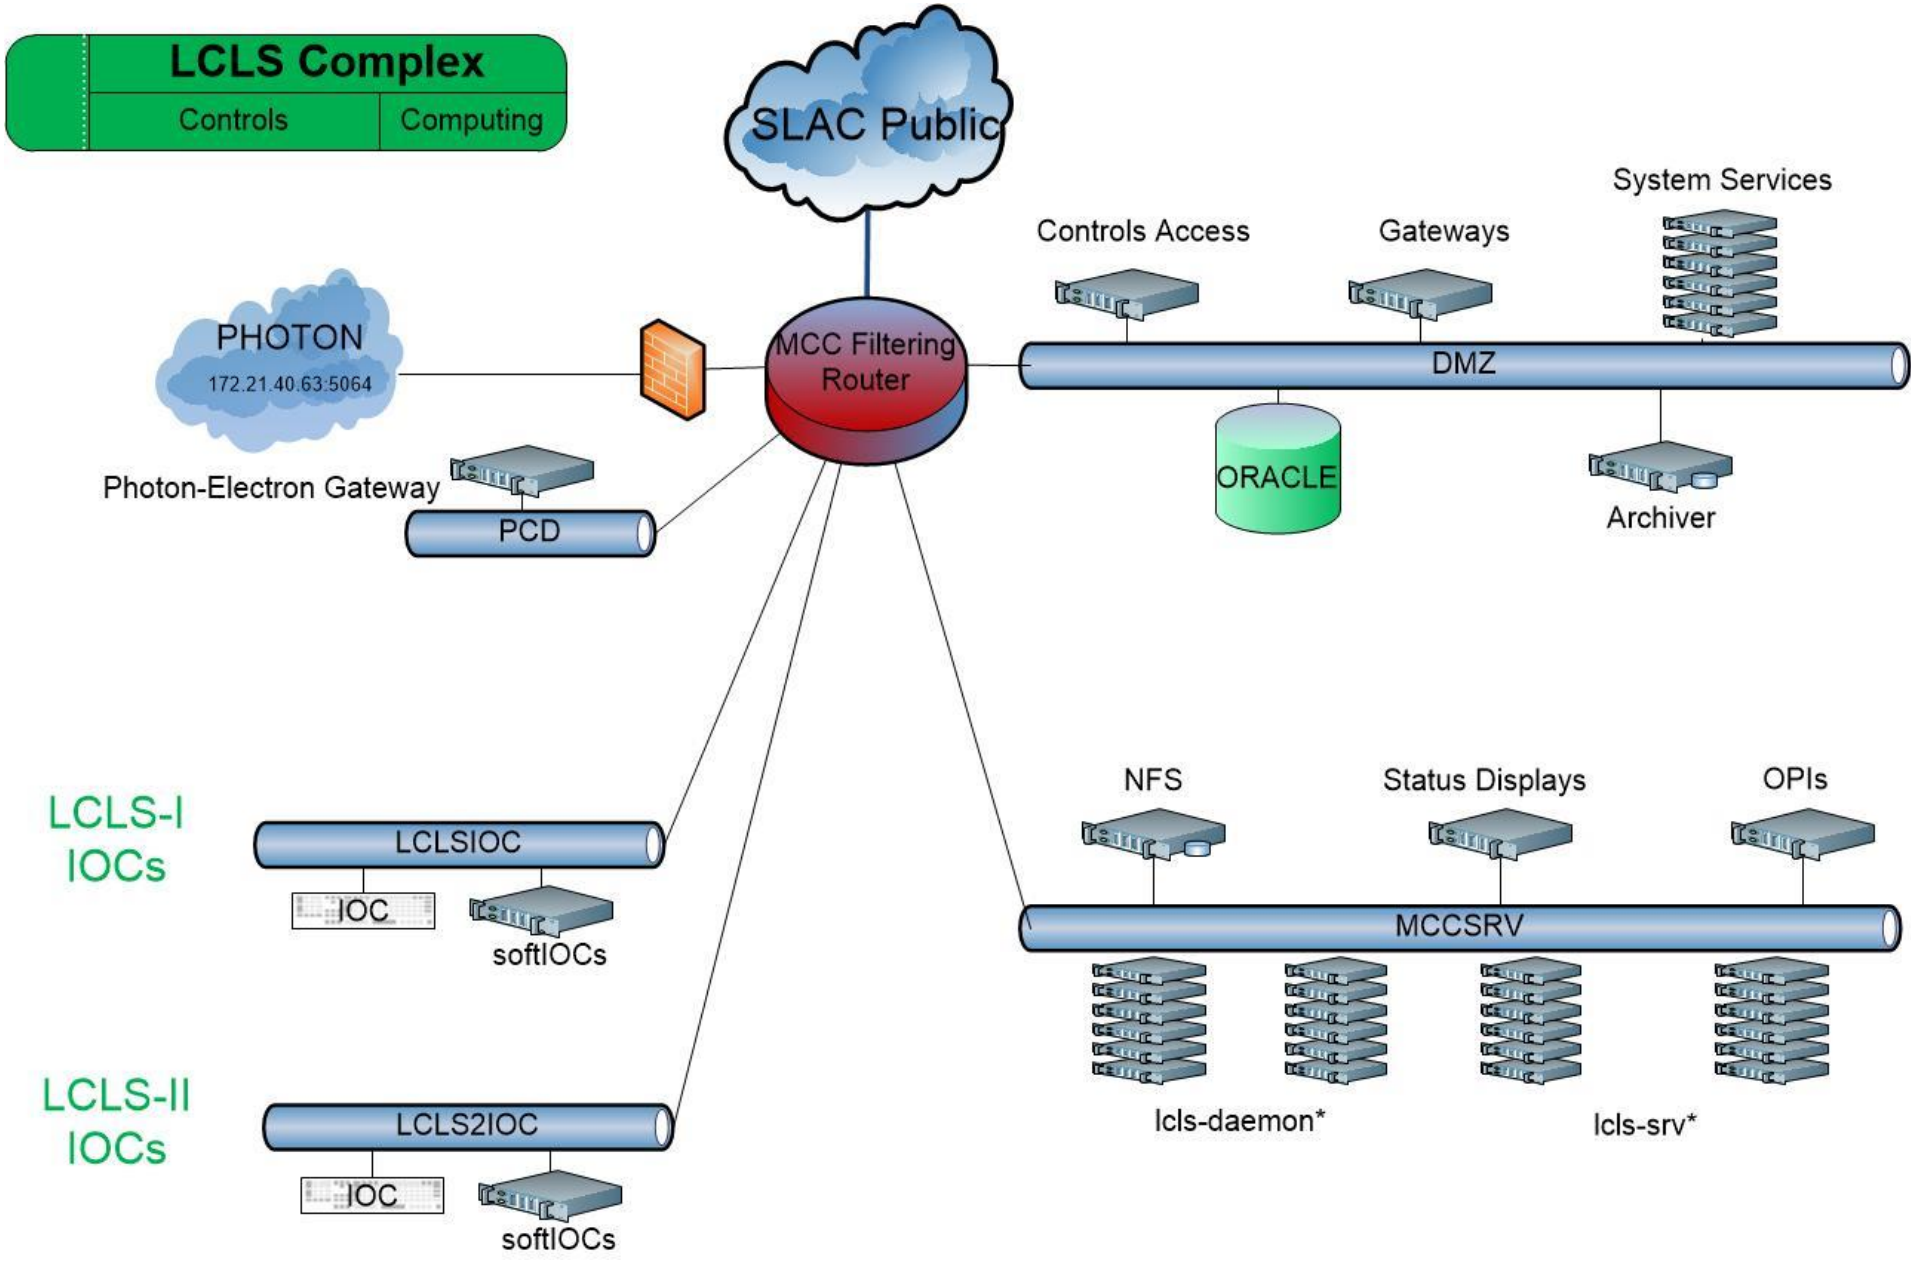

Jingchen Zhou

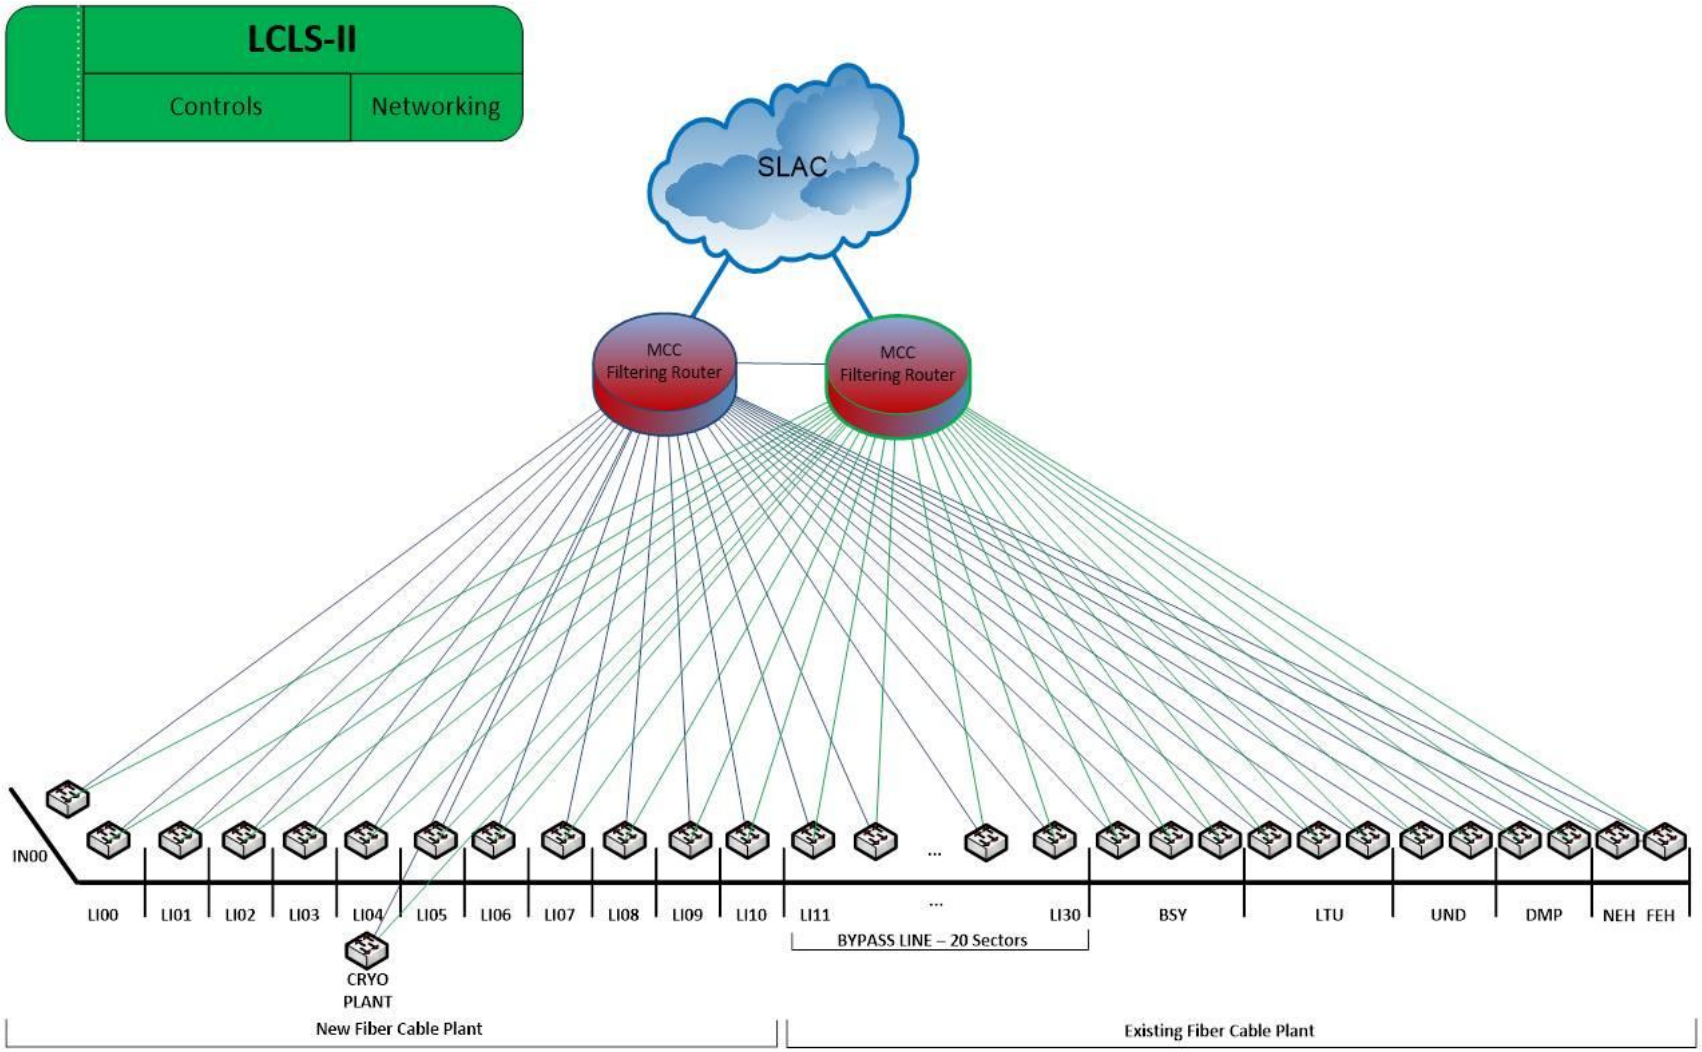

- **Network has redundant Core routers in MCC and redundant fiber links to all switches**

Charles Granieri

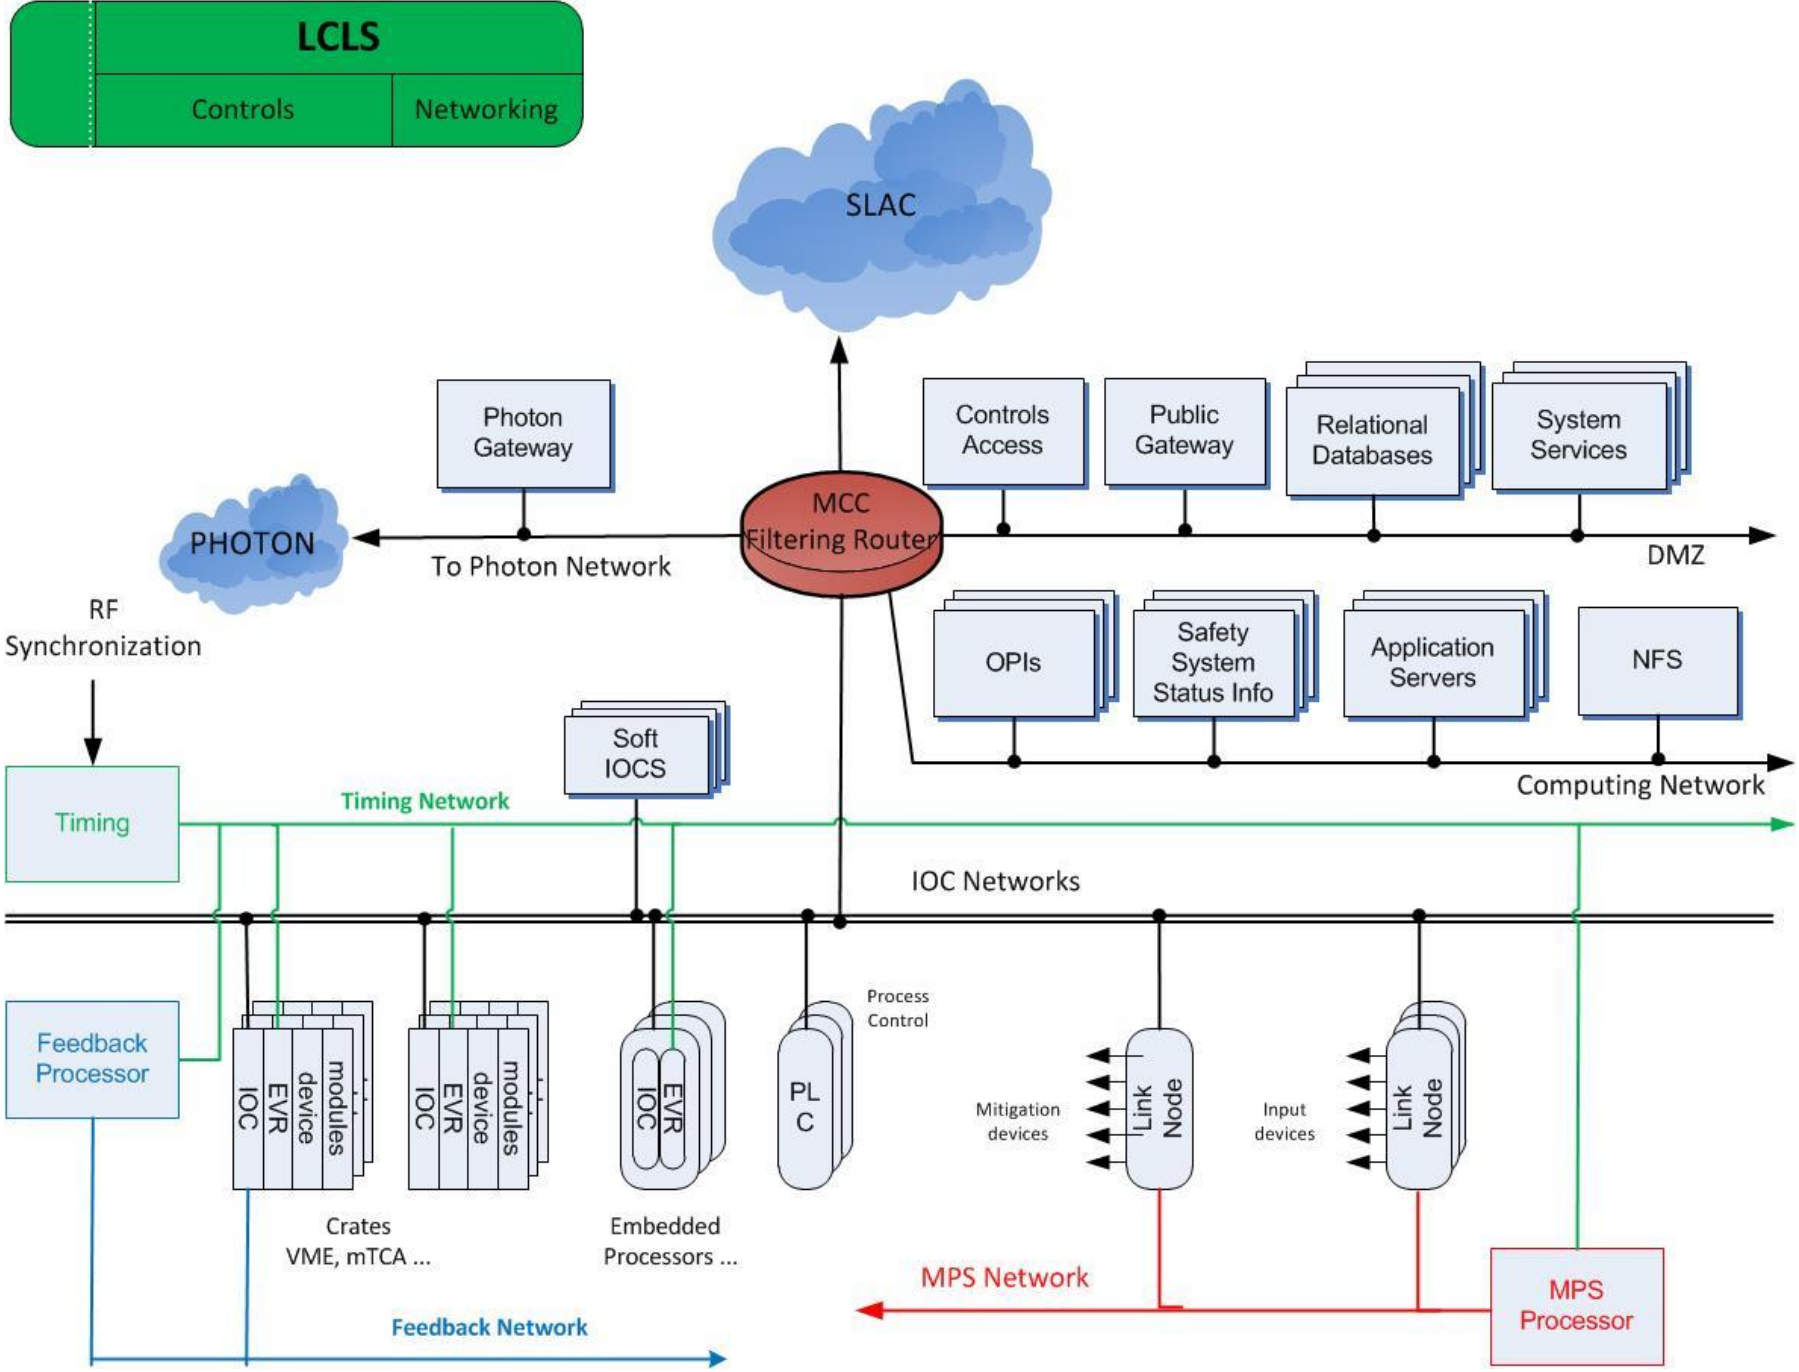
LCLS
Controls
Networking
SLAC
Photon Gateway
Controls Access
Public Gateway
Relational Databases
System Services
PHOTON
To Photon Network
MCC Filtering Router
DMZ
RF Synchronization
OPIs
Safety System Status Info
Application Servers
NFS
Soft IOCS
Timing
Timing Network
Computing Network
IOC Networks
Feedback Processor
IOC
EVR
device
modules
Process Control
PLC
Mitigation devices
Link Node
Input devices
Crates
VME, mTCA ...
Embedded Processors ...
MPS Network
MPS Processor
Feedback Network

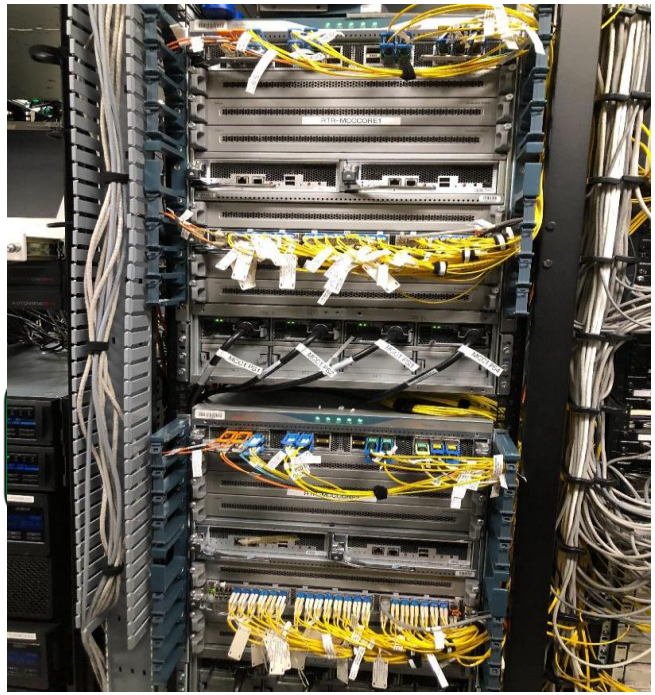

MCC Routers

Accelerator Area

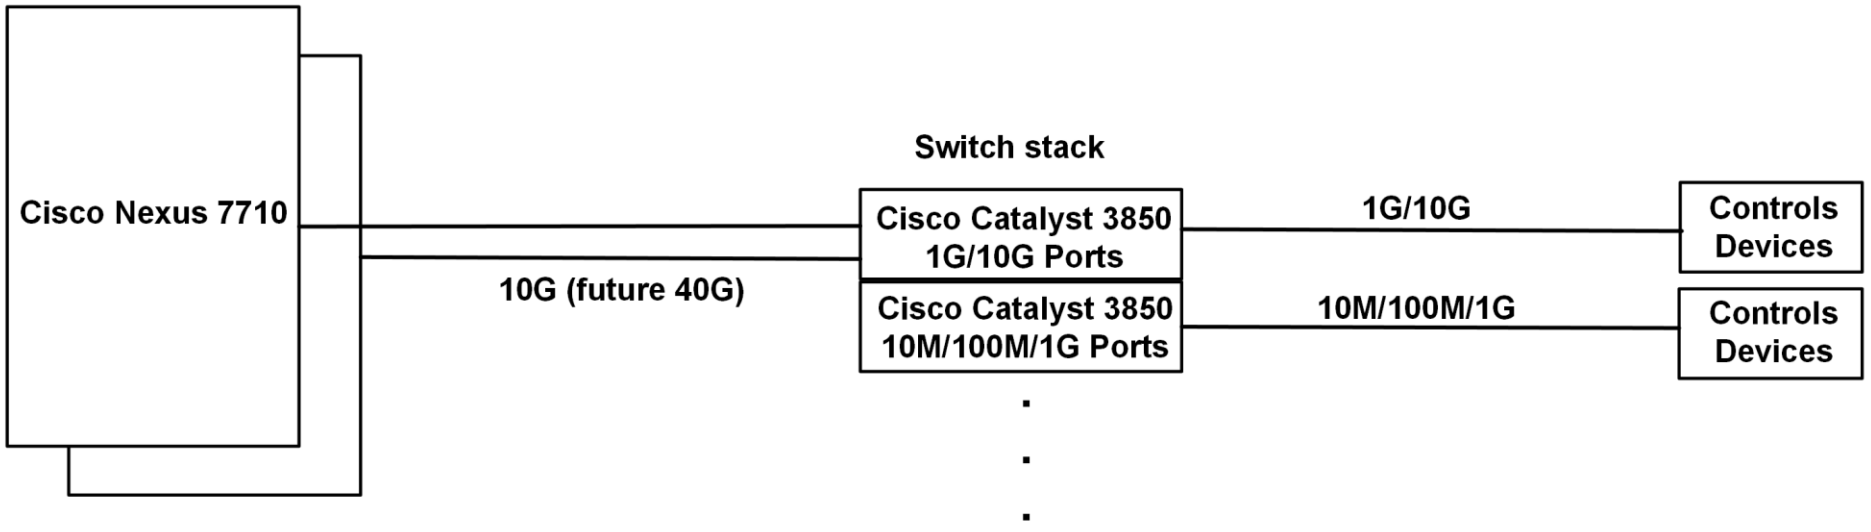

Charles Granieri

NFS Server: Oracle ZFS Storage System (ZS7-2)

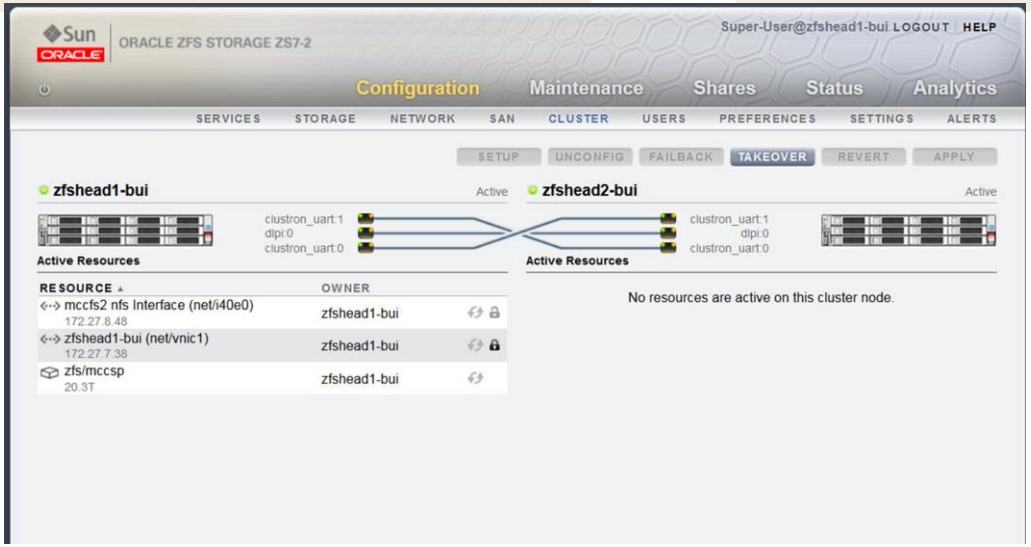

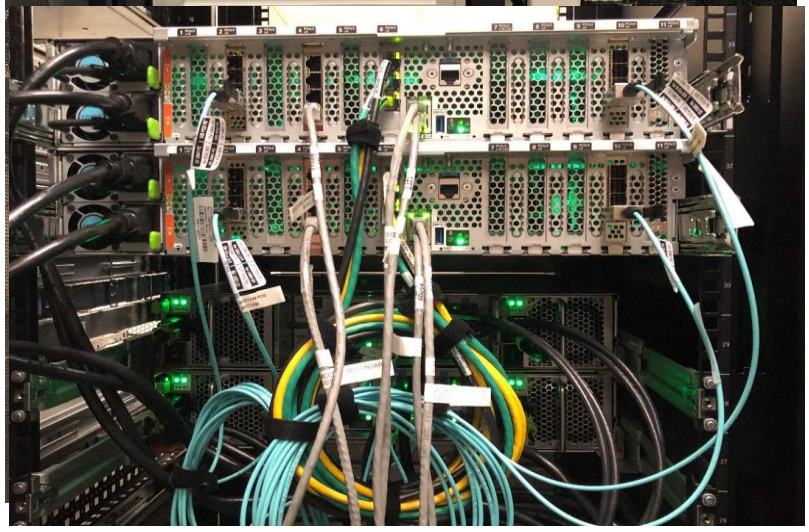

- Significant improvement in performance, particularly for controls applications that require a lot of dynamic libraries
- Highly reliable, as the most critical system in Controls Computing Infrastructure
- Compatible with Controls Backup/Restore System via NDMP protocol
- Compatible with NFSv2 required by legacy RTEMS IOCs.
- Enhanced DTrace Analytics for real-time analysis and monitoring

Application Performance Study

SLAC

Loading Shared Libraries

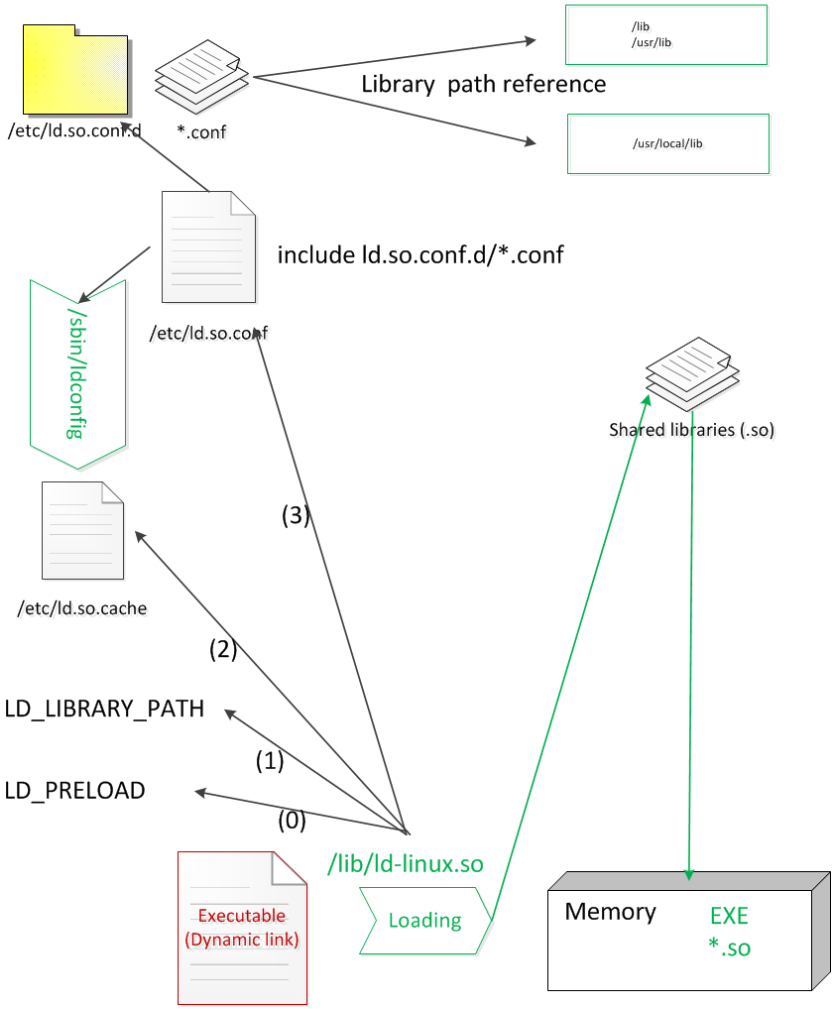

Jingchen Zhou

Disk Caching

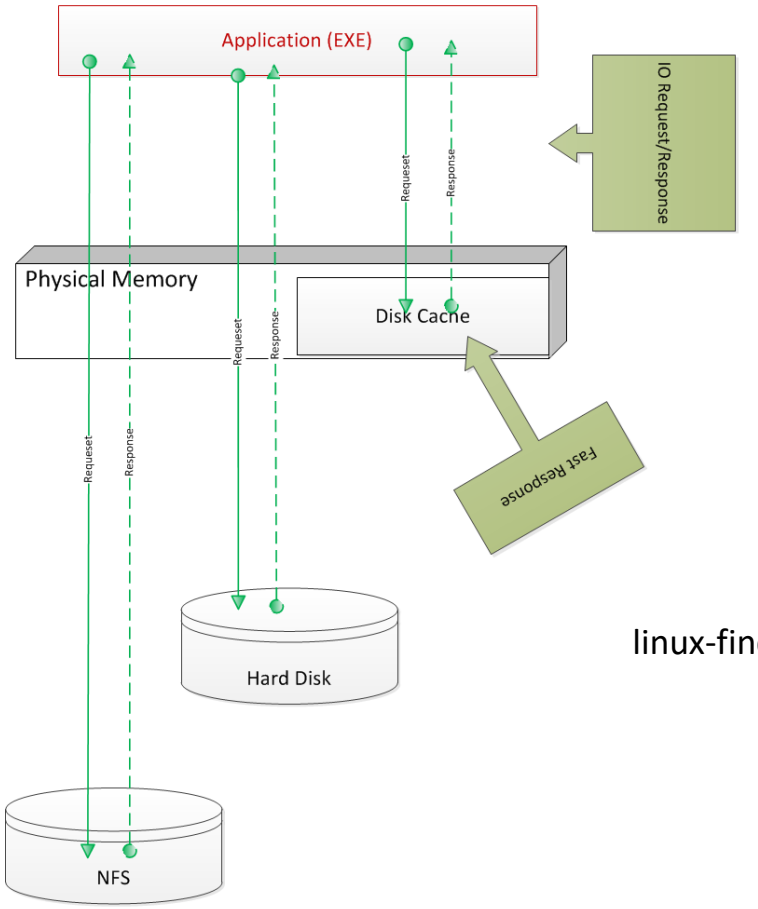

linux-fincore

- Cold start
  Loaded from hard disk if locally or from NFS if distributed
- Warm start
  Loaded from Disk Cache in Memory

All in memory while running.

DTrace Analytics

- Provides real-time analysis and monitoring functionality, enabling fine-grained visibility into disk, flash, CPU, networking, data service (NFS), and other statistics
- Supports rapid identification and resolution of bottlenecks for troubleshooting and performance tuning

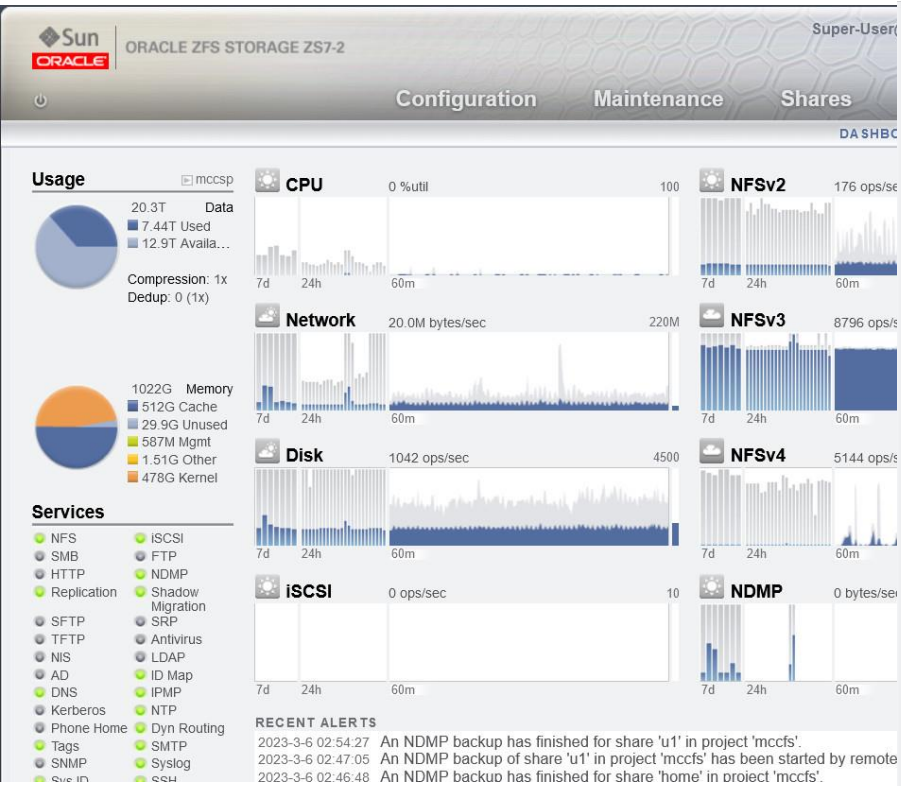

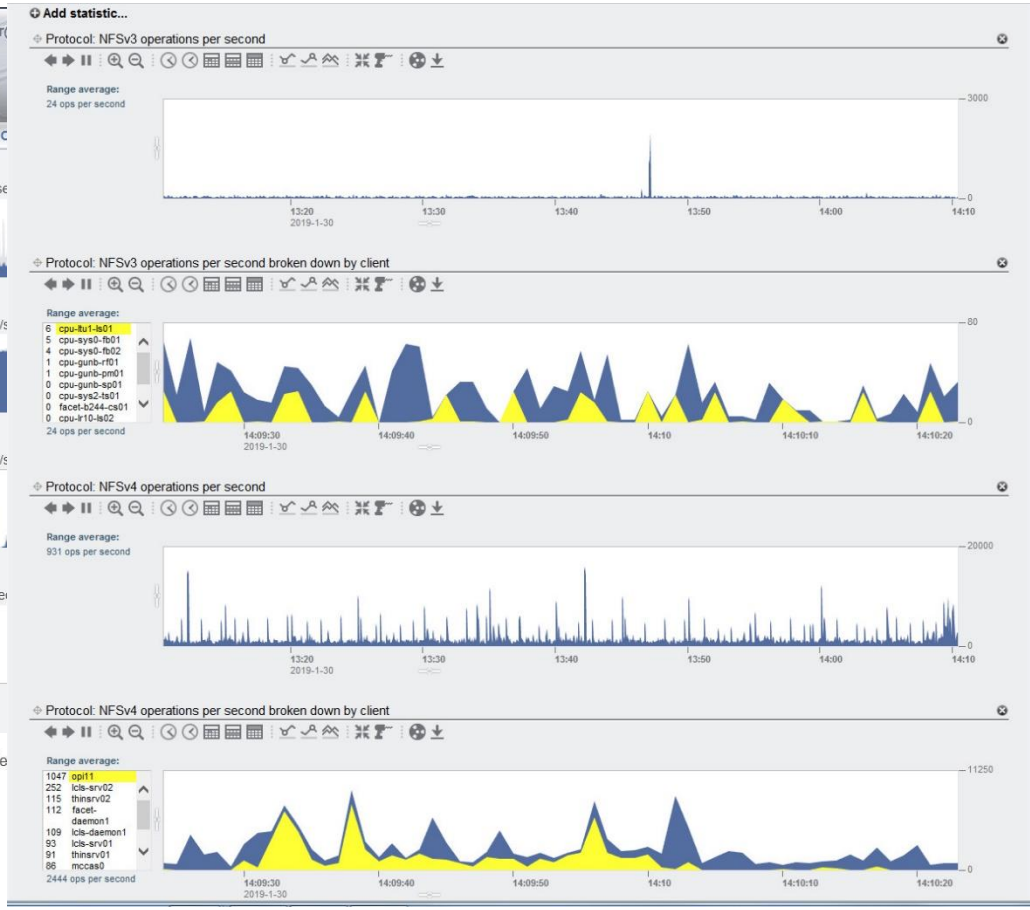

Network Service (DHCP, PXE, TFTP, NTP, etc.)

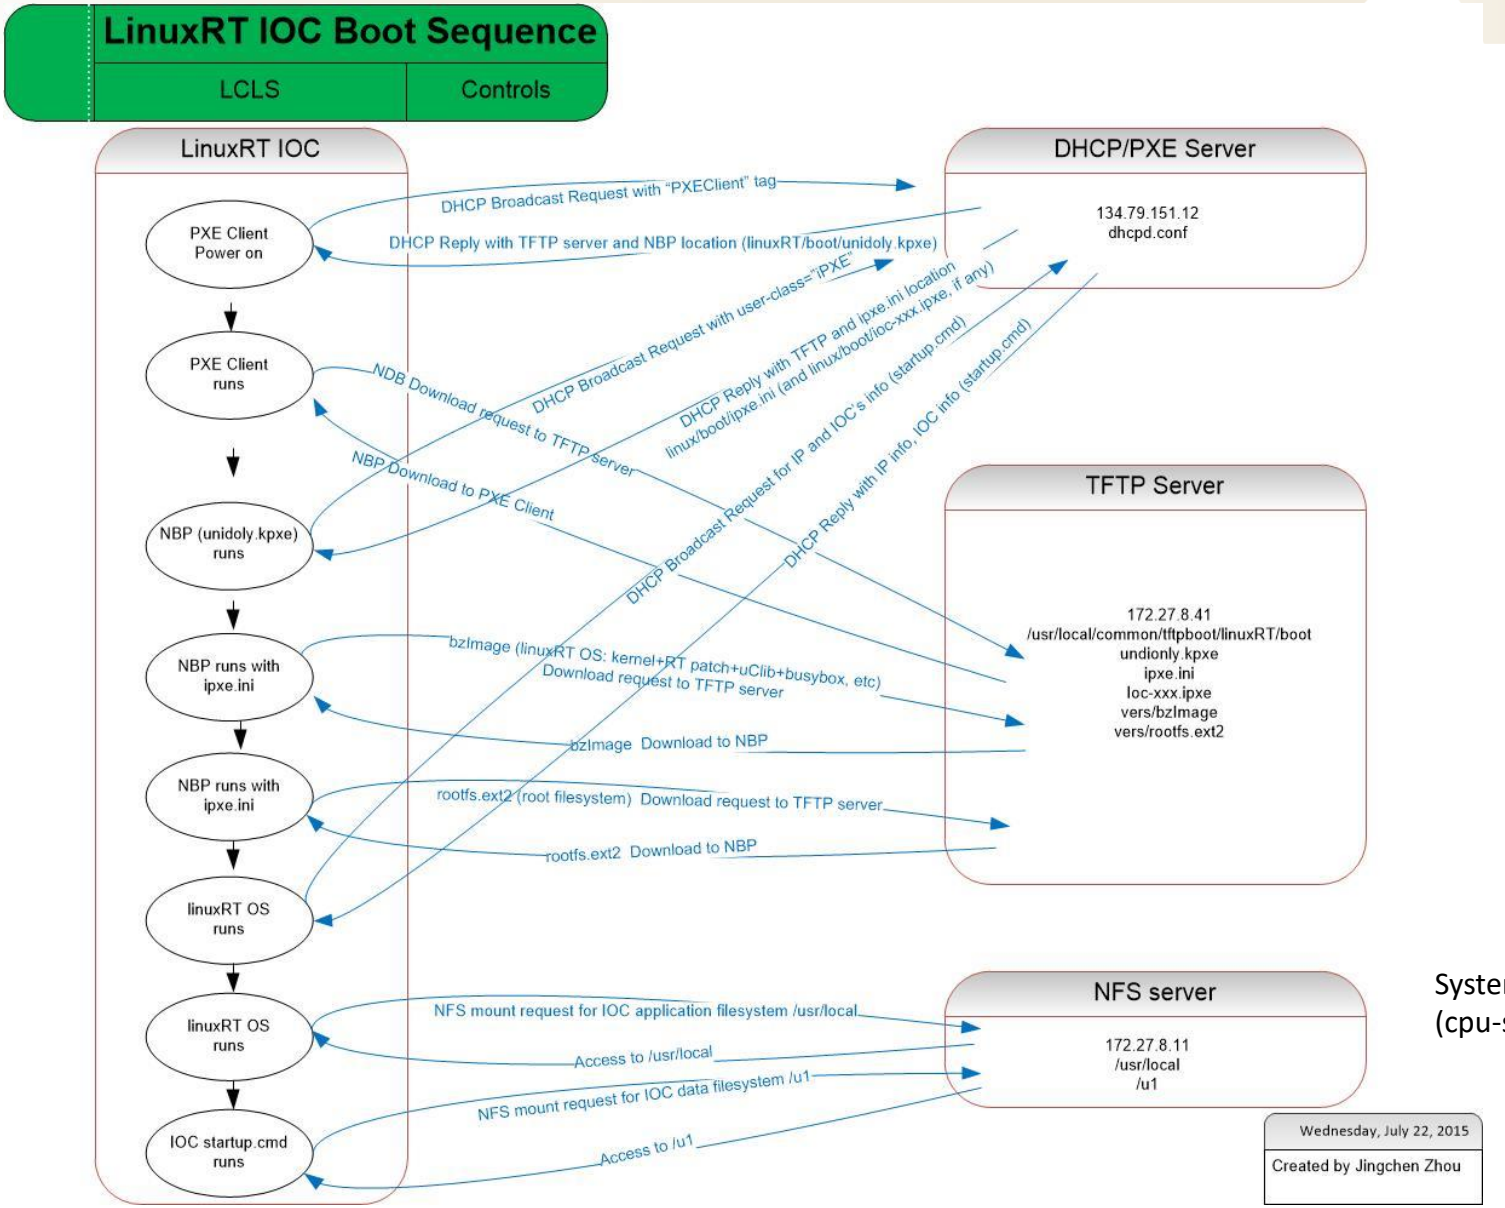

System Test linuxRT IOC (cpu-sys0-ck00)

Thin Client System

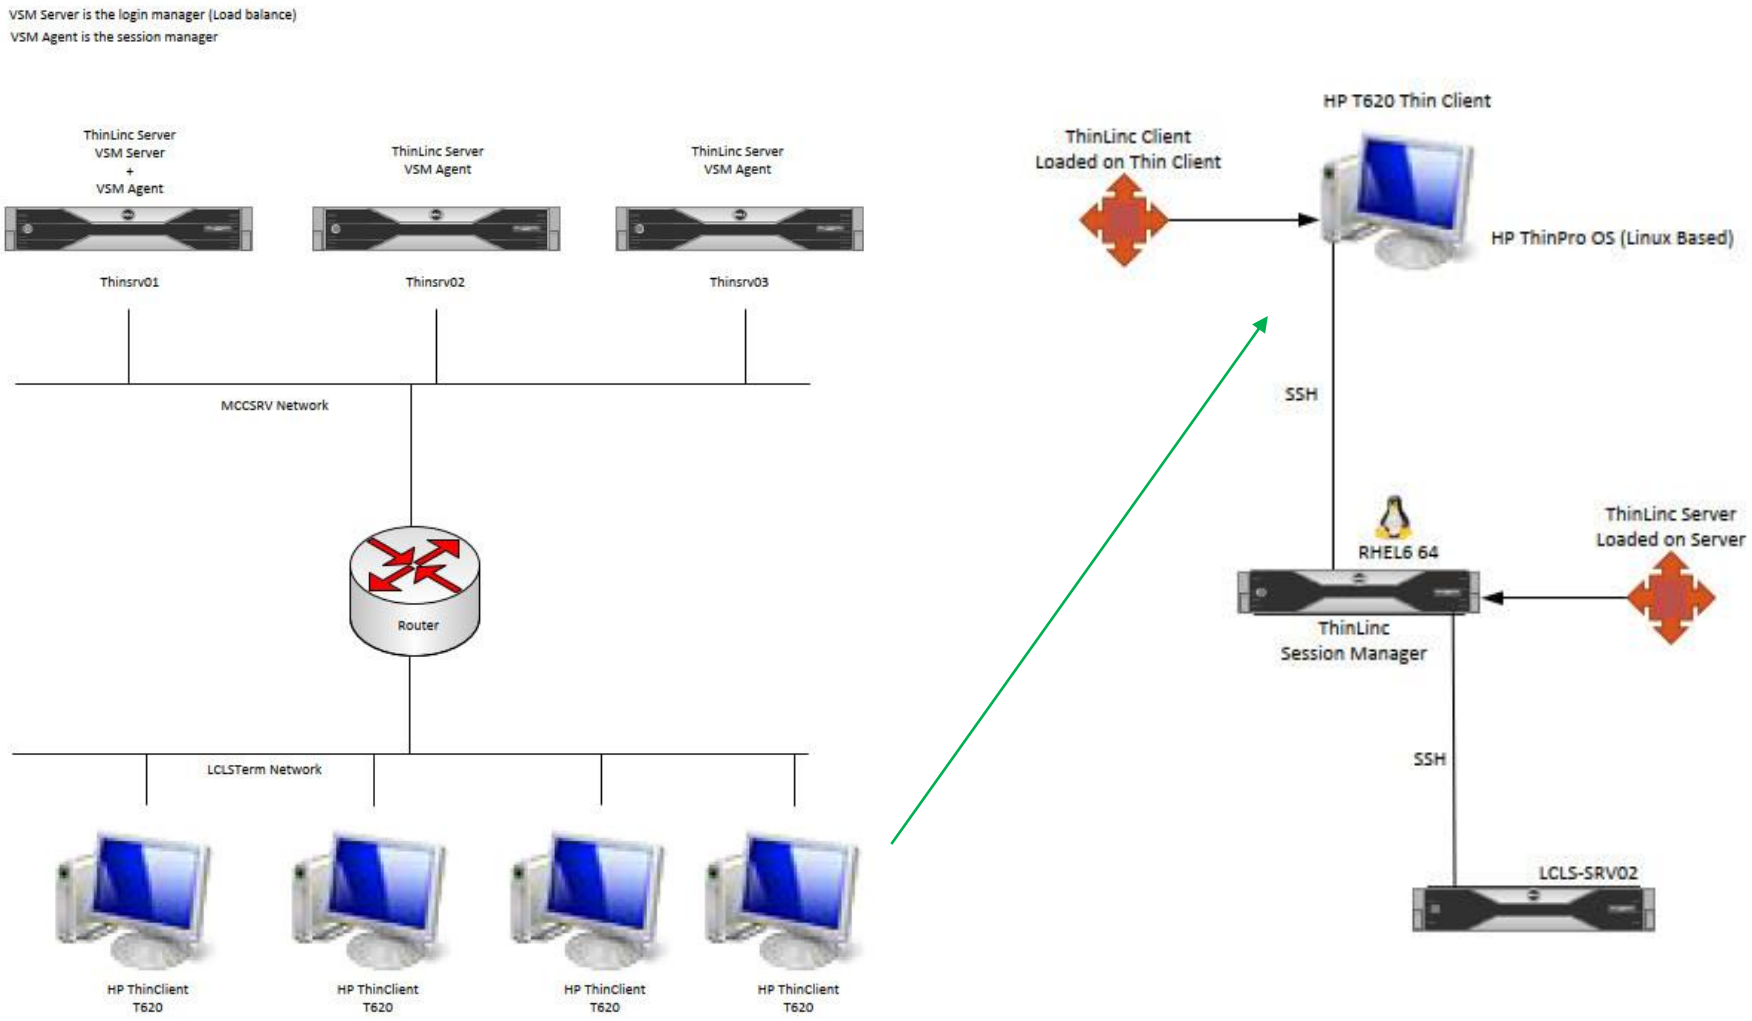

Backup/Restore System

- Controls Backup/Restore System
  - EMC Data Domain
  - Networker software

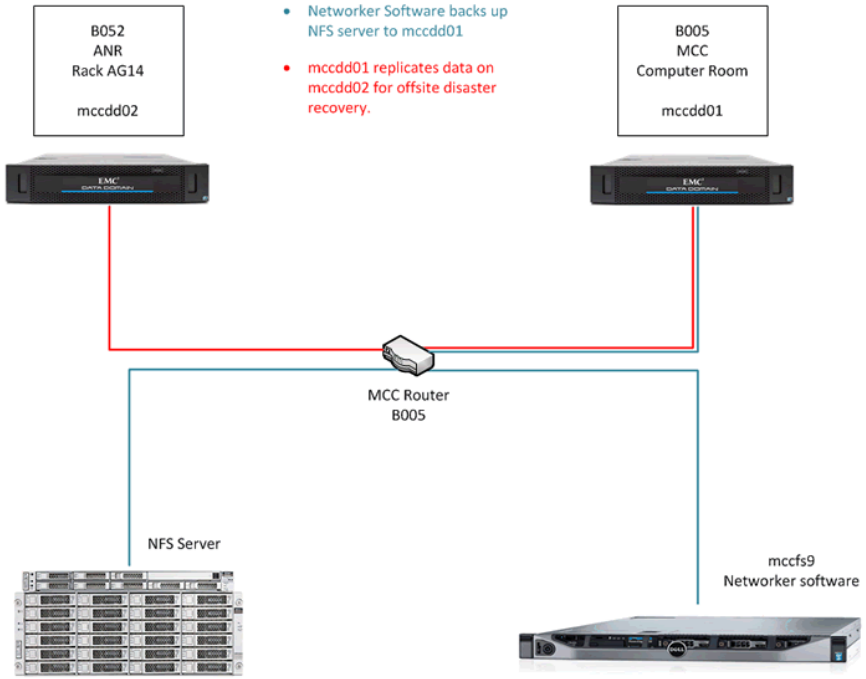

Ken Brobeck

- NDMP (Network Data Management Protocol), a protocol to optimize Backup performance for NFS based data transportation

ACR and ANR

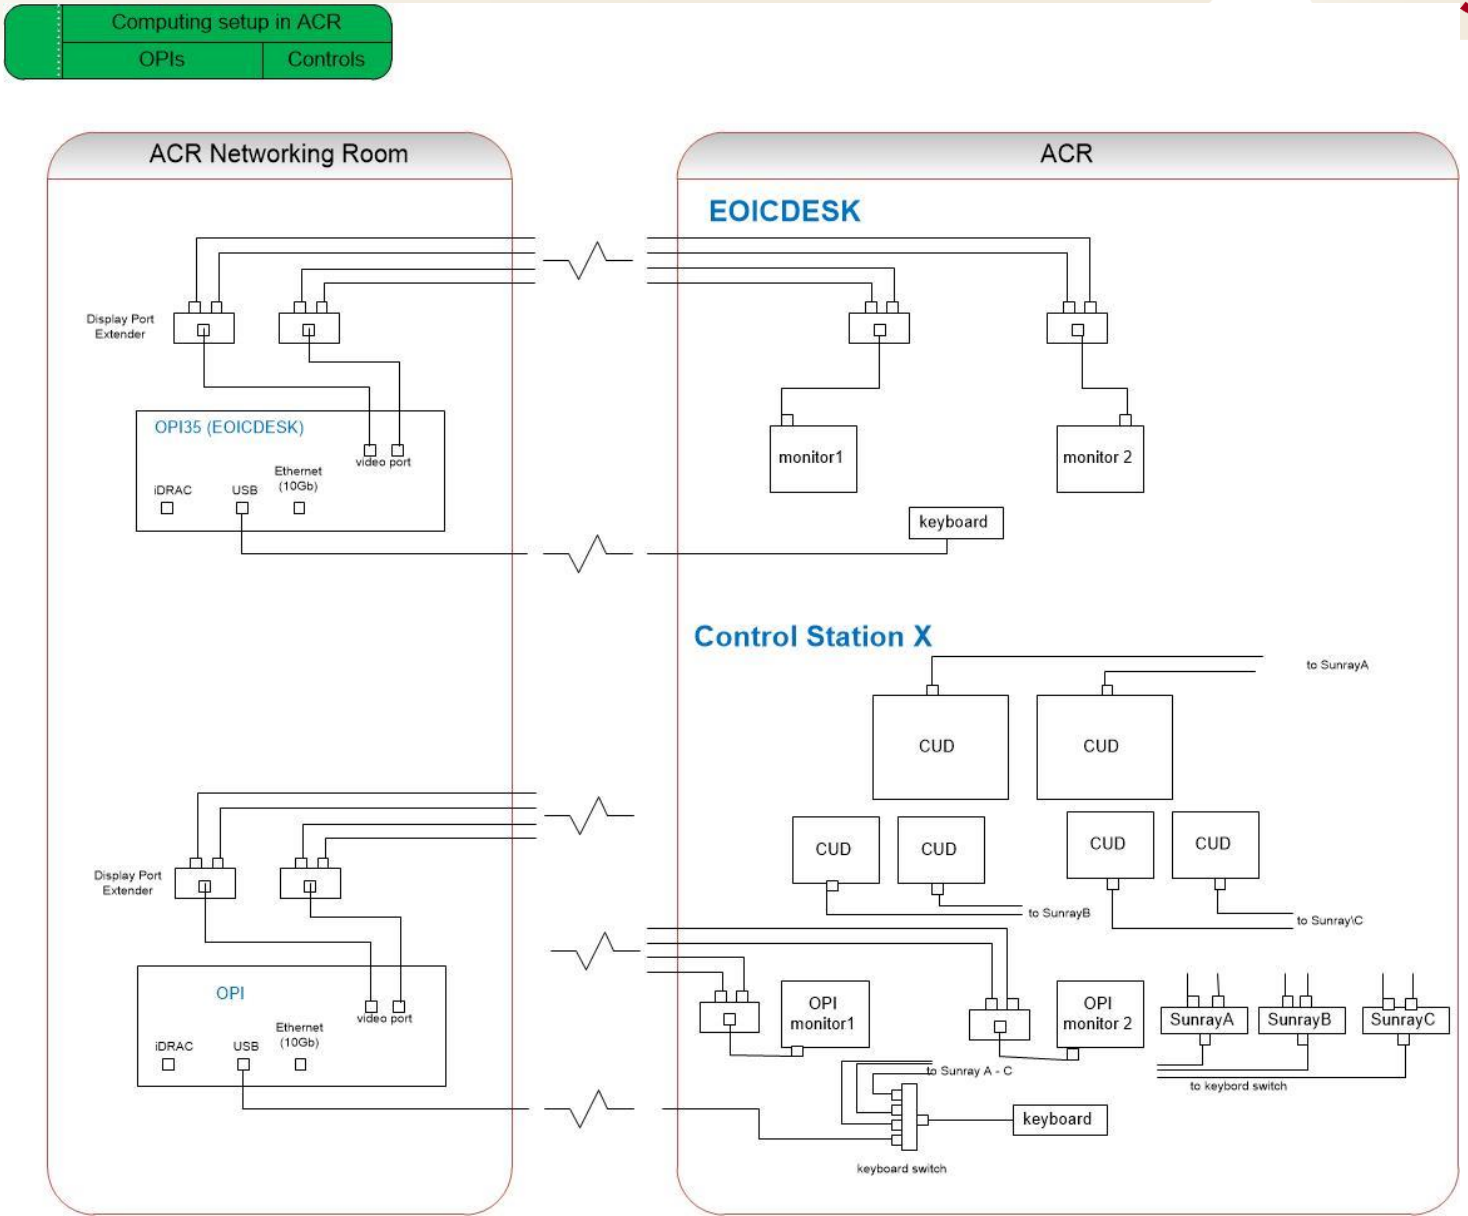
Computing setup in ACR
OPIs
Controls
ACR Networking Room
ACR
EOICDESK
Display Port Extender
OPI35 (EOICDESK)
iDRAC
USB
Ethernet (10Gb)
video port
monitor1
monitor 2
keyboard
Control Station X
to SunrayA
CUD
CUD
CUD
CUD
CUD
CUD
to SunrayB
to SunrayC
Display Port Extender
OPI
iDRAC
USB
Ethernet (10Gb)
video port
OPI monitor1
OPI monitor 2
SunrayA
SunrayB
SunrayC
to Sunray A - C
to keybord switch
keyboard
keyboard switch

Controls PV Gateways

Name	Purpose	PVs served	Host	CAS port	Beacon Port	Beacon sent to	Document/Comment
LCLS							
gwEbeamServe	Serving LCLS PVs (Ebeam) to Photon	Readonly Write for selected PVs	lcls-daemon3	5080	5081	172.21.40.63 (Photon Gateway Subnet)	Design Admin Allow all Deny OTRS:DMP1:695:.*
gwEbeamServeWF	Serving LCLS PVs (Ebeam) to Photon	Readonly Write for selected PVs	lcls-daemon3	5079	5081	172.21.40.63 (Photon Gateway Subnet)	Deny all Allow OTRS:DMP1:695:.* (dedicated to serve XTCAV OTRDMP camera image PV)
gwLCLS4FACET	Serving LCLS PVs to FACET	Readonly	lcls-daemon10	5070	5069	172.27.75.255 (FACETCA)	Design Admin
gwLCLSPUB	Serving LCLS PVs to public	Readonly	lcls-prod01	5068	5069	134.79.151.255 (DMZ)	Admin
gwLCLSARCH0	Serving LCLS and LCLS-II PVs to Archiver on DMZ	Readonly	lcls-prod01	5076	5069	134.79.151.255 (DMZ)	.* ALLOW (default) Deny a list refer gwLCLSARCH*.dat
gwLCLSARCH1	Serving LCLS and LCLS-II PVs to Archiver on DMZ	Readonly	lcls-prod01	5077	5069	134.79.151.255 (DMZ)	.* DENY Allow portion in the list refer gwLCLSARCH*.dat
gwLCLSARCH2	Serving LCLS and LCLS-II PVs to Archiver on DMZ	Readonly	lcls-prod01	5078	5069	134.79.151.255 (DMZ)	.* DENY Allow portion in the list refer gwLCLSARCH*.dat
gwLCLSARCH3	Serving LCLS and LCLS-II PVs to Archiver on DMZ	Readonly	lcls-prod01	5079	5069	134.79.151.255 (DMZ)	.* DENY Allow portion in the list refer gwLCLSARCH*.dat
gwLCLSARCH4	Serving LCLS and LCLS-II PVs to Archiver on DMZ	Readonly	lcls-prod01	5074	5069	134.79.151.255 (DMZ)	.* DENY Allow portion in the list refer gwLCLSARCH*.dat
gwLCLSARCH5	Serving LCLS and LCLS-II PVs to Archiver on DMZ	Readonly	lcls-prod01	5075	5069	134.79.151.255 (DMZ)	.* DENY Allow portion in the list refer gwLCLSARCH*.dat
gwLCLSARCH6	Serving LCLS and LCLS-II PVs to Archiver on DMZ	Readonly	lcls-prod01	5072	5069	134.79.151.255 (DMZ)	.* DENY Allow portion in the list refer gwLCLSARCH*.dat
gwLCLSARCH7	Serving LCLS and LCLS-II PVs to Archiver on DMZ	Readonly	lcls-prod01	5073	5069	134.79.151.255 (DMZ)	.* DENY Allow portion in the list refer gwLCLSARCH*.dat
gwCRYO4LCLS	Serving Cryoplant PVs to LCLS-II	Read and Write	cryo-daemon1	5061	5069	172.27.43.255 (CRYO) 172.27.11.255 (MCCSRV) 134.79.151.21 (LCLS-PROD01)	
FACET							
gwFACET4LCLS	Serving FACET PVs to LCLS	Readonly	facet-daemon1	5070	5069	172.27.11.255 (LCLSCA)	Design Admin
gwFACETPUB	Serving FACET PVs to public	Readonly	lcls-prod01	5063	5069	134.79.151.255 (DMZ)	Design
gwFACETARCH0	Serving FACET PVs to Archiver on DMZ	Readonly	lcls-prod01	5064	5069	134.79.151.255 (DMZ)	.* ALLOW (default) Deny a list refer gwFACETARCH*.dat
gwFACETARCH1	Serving FACET PVs to Archiver on DMZ	Readonly	lcls-prod01	5062	5069	134.79.151.255 (DMZ)	.* DENY Allow portion in the list refer gwFACETARCH*.dat
Test Facilities							
gwACCTESTPUB	Serving Test Facilities PV to public	Readonly	testfac-daemon2	5048	5049	134.79.219.255 (LCLSDEV)	doc

LCLS Electron-Photon Gateways:
https://www.slac.stanford.edu/grp/cd/soft/unix/slaconly/ElectronPhotonGateway.html

LCLS Archiver System

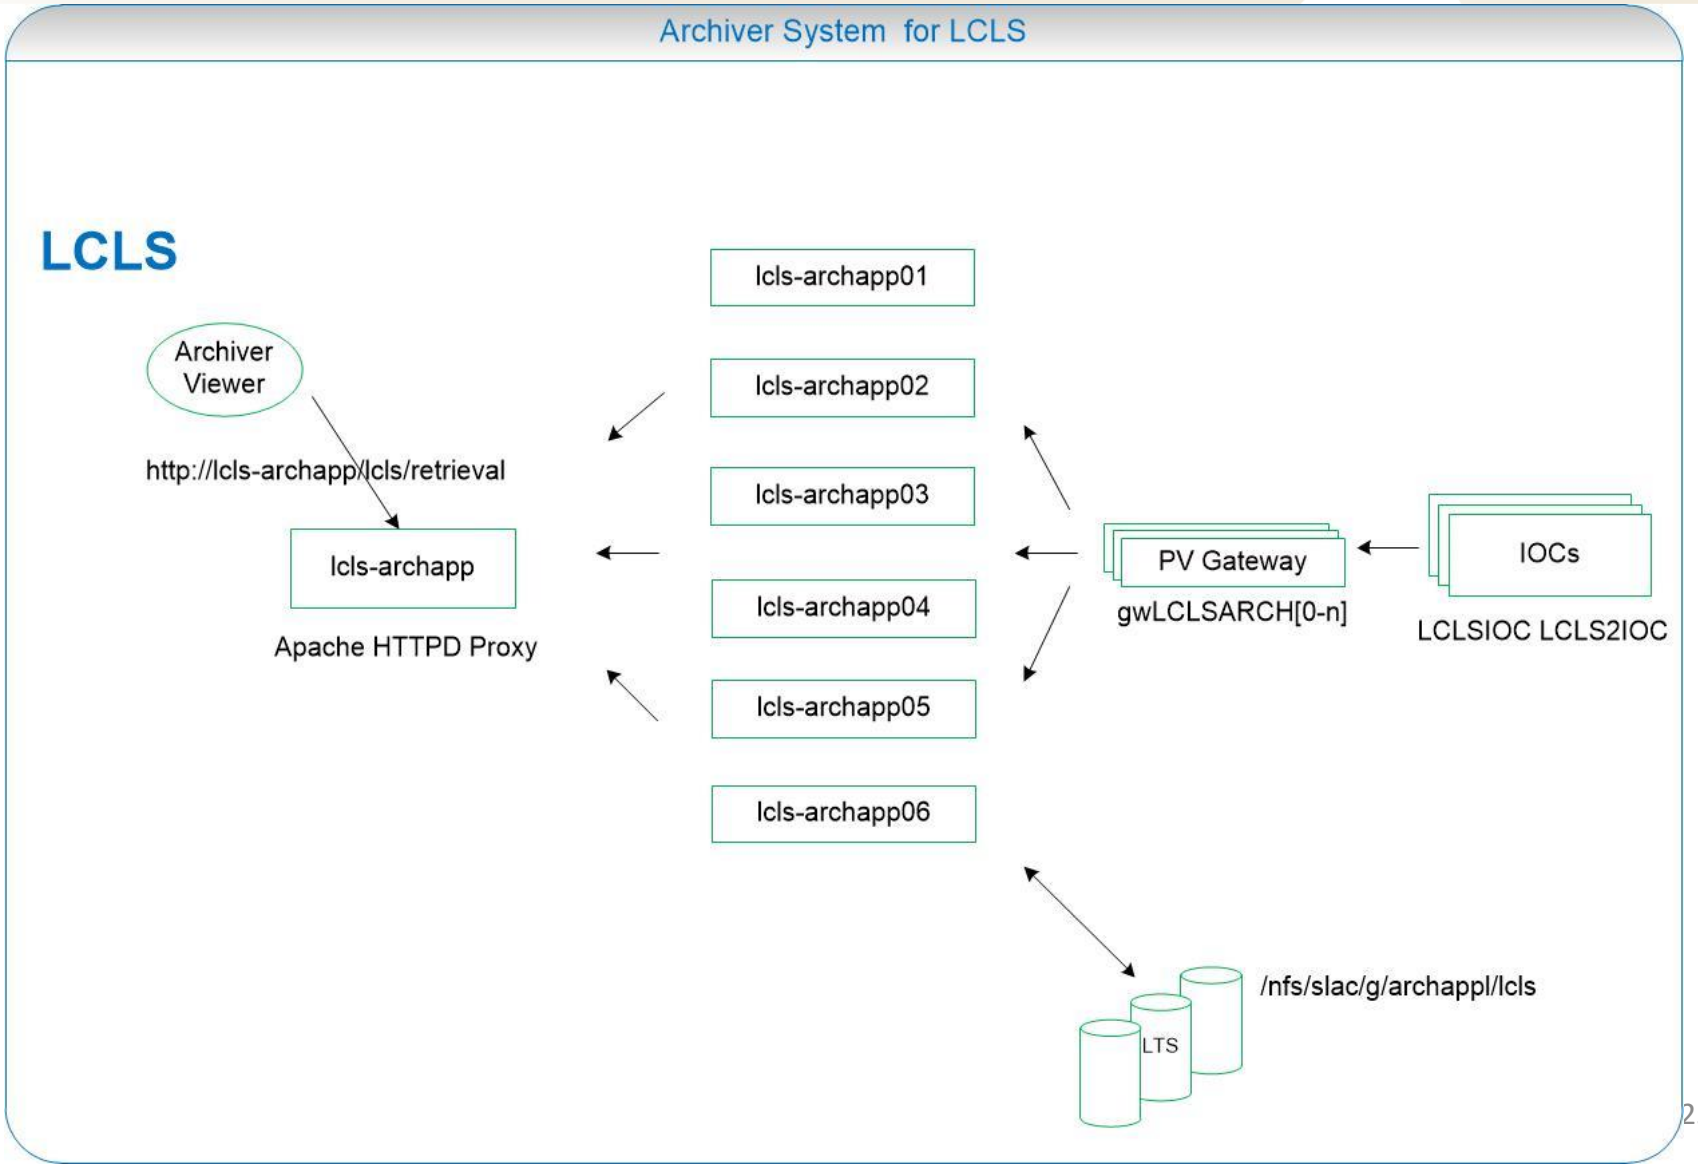

Controls Network Monitoring

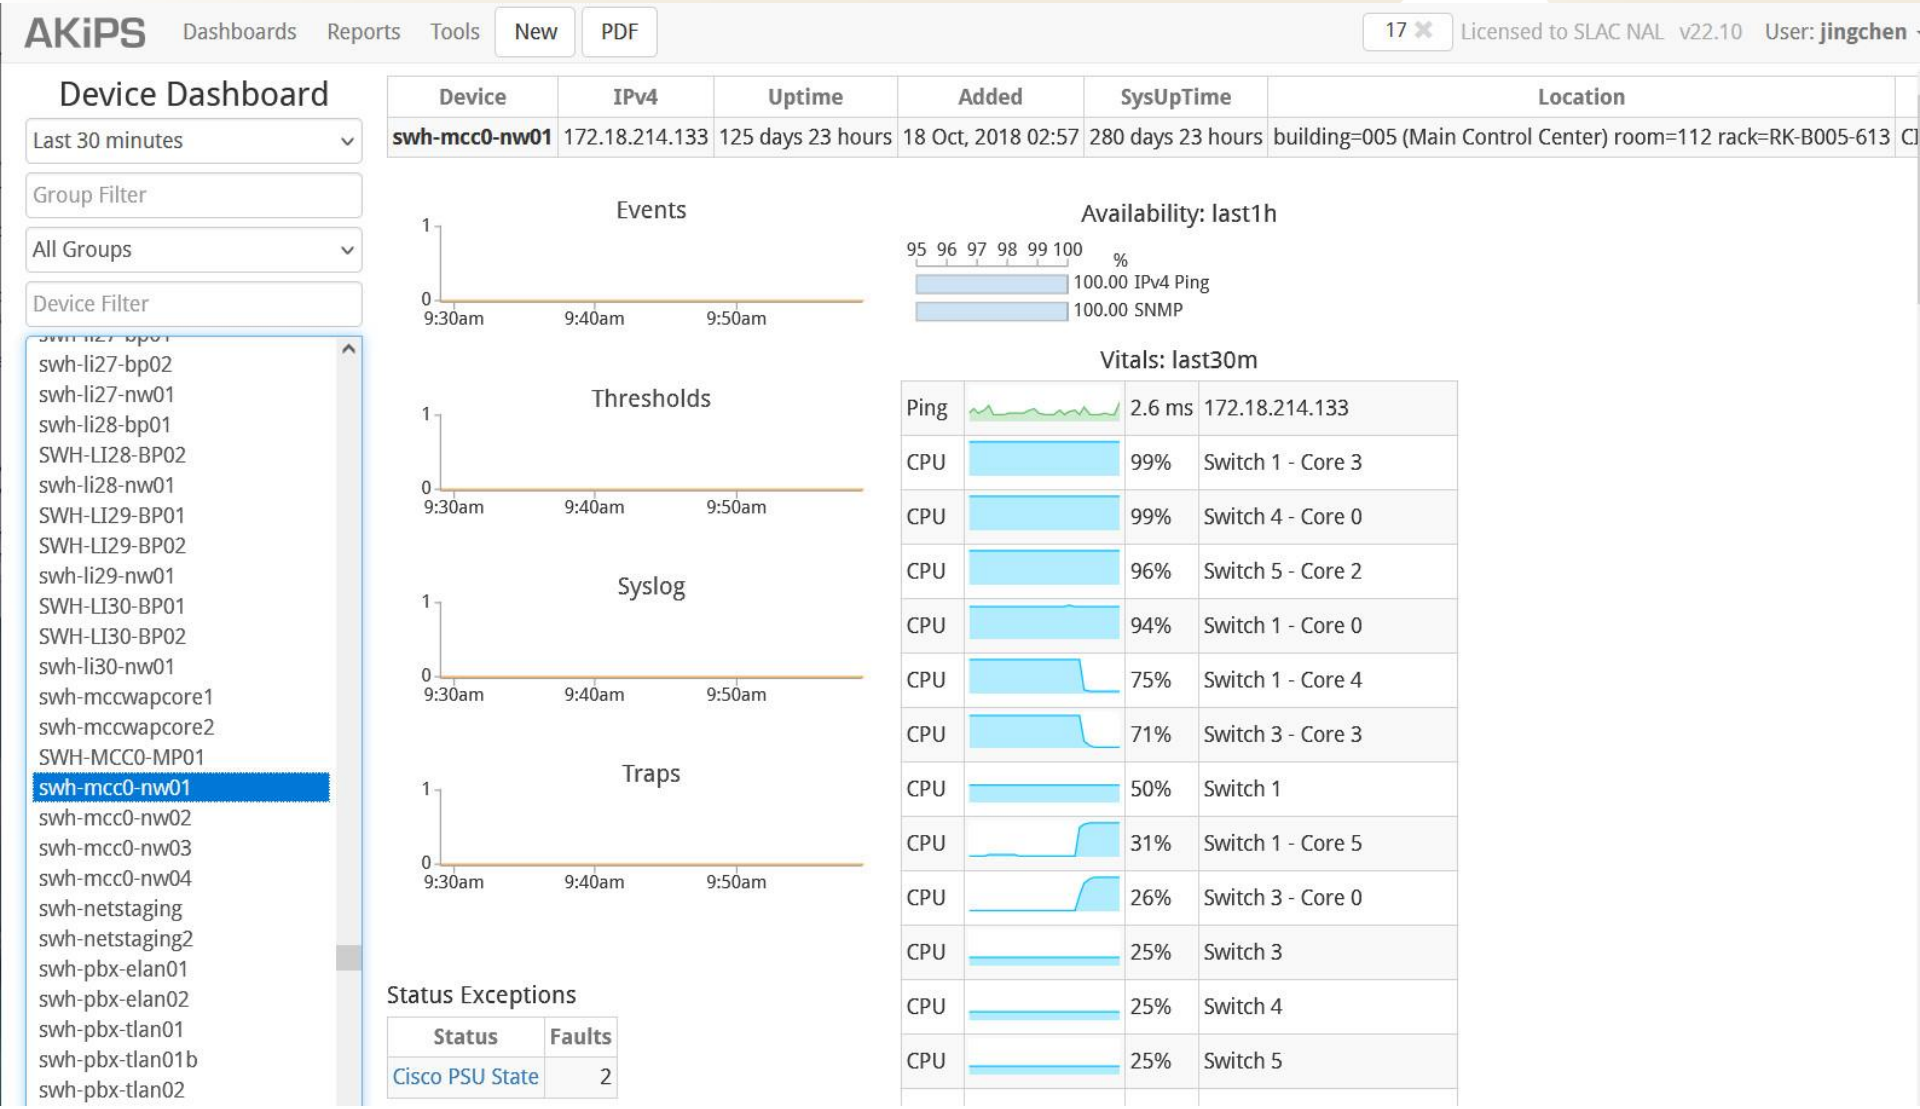

Controls Network Monitoring

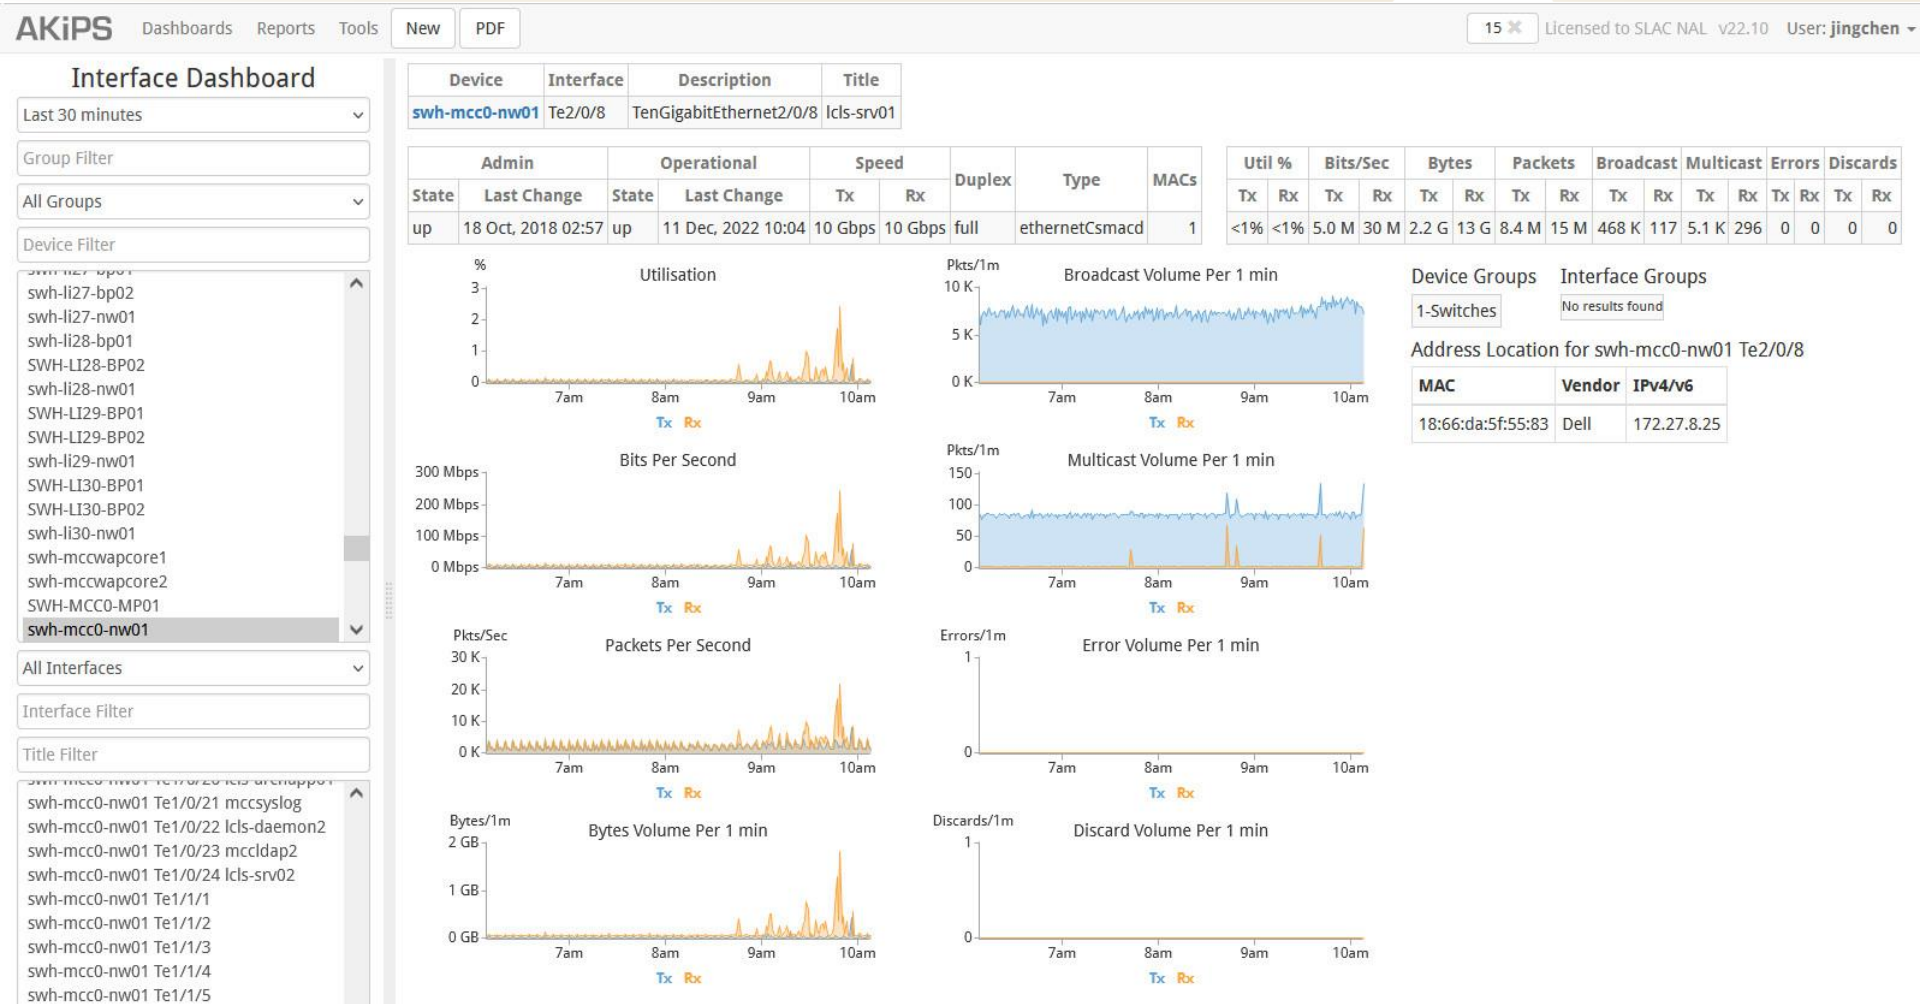

Controls Network Troubleshooting

- Network troubleshooting
  - Network troubleshooting server: lcls-srv05
  - Core router and lcls-srv05 configured to monitor and troubleshoot controls subnets, e.g.,
    - p1p1 : to monitor and troubleshoot network packets on LCLSUTIL/LCLS2UTIL/MCCSRV
    - p1p2 : to monitor and troubleshoot network packets on FNET/LCLSIOC/LCLS2IOC
- Network tools
  - tcpdump
  - wireshark
  - lsof, iftop, nethogs, ttcp, etc.
- Latest example: Controls network overloaded by PVA multicast searches due to a bug in java implementations of EPICSv7 PVA

Reliability and Security

- Reliability:
  - Critical systems with full redundancy (NFS, DHCP, DNS, NTP, Oracle, etc.)
  - All other systems with a system mirroring and a failover procedure
  - A Backup/Restore system for all data and software
- Security:
  - Cyber Security Program Plan
    - MCC Enclave Security Controls
    - A supplement to SLAC Core Security Program Plan
    - Reviewed by DOE
  - Protected networks for all critical computing systems
    - Limited access only provided to authorized SLAC public users via a dedicated Login system on DMZ
    - Read-only PVs provided to SLAC public via a PV gateway on DMZ
  - CrowdStrike applied to systems visible to SLAC public.
  - Any vulnerability issues fixed timely
  - DOE penetration test coming soon

The detailed efforts behind the scenes

On-going efforts required for sustaining engineering and maintenance to provide networking and computing infrastructures support for LCLS, LCLS-II, FACET-II, Test Facilities, and various R&D efforts for SLAC accelerator controls.

- Computing Infrastructure Support for Accelerator Controls

https://www.slac.stanford.edu/grp/cd/soft/unix/slaconly/computing_support.pdf

- Networking Infrastructure Support for Accelerator Controls

https://www.slac.stanford.edu/grp/cd/soft/unix/slaconly/networking_support.pdf

Summary

- Dedicated to providing the infrastructures highly critical to SLAC accelerator operations for
  - Test Facilities
  - FACET-II
  - LCLS and LCLS-II
- Skillful in tackling problems with complexities
- Striving to deliver sustainable solutions to meet growing demands
- Responsible for networking and computing infrastructure design for accelerator controls
- The listed on-going efforts required for sustaining engineering and maintenance
- The infrastructures: secure and reliable
- Modernization is coming