

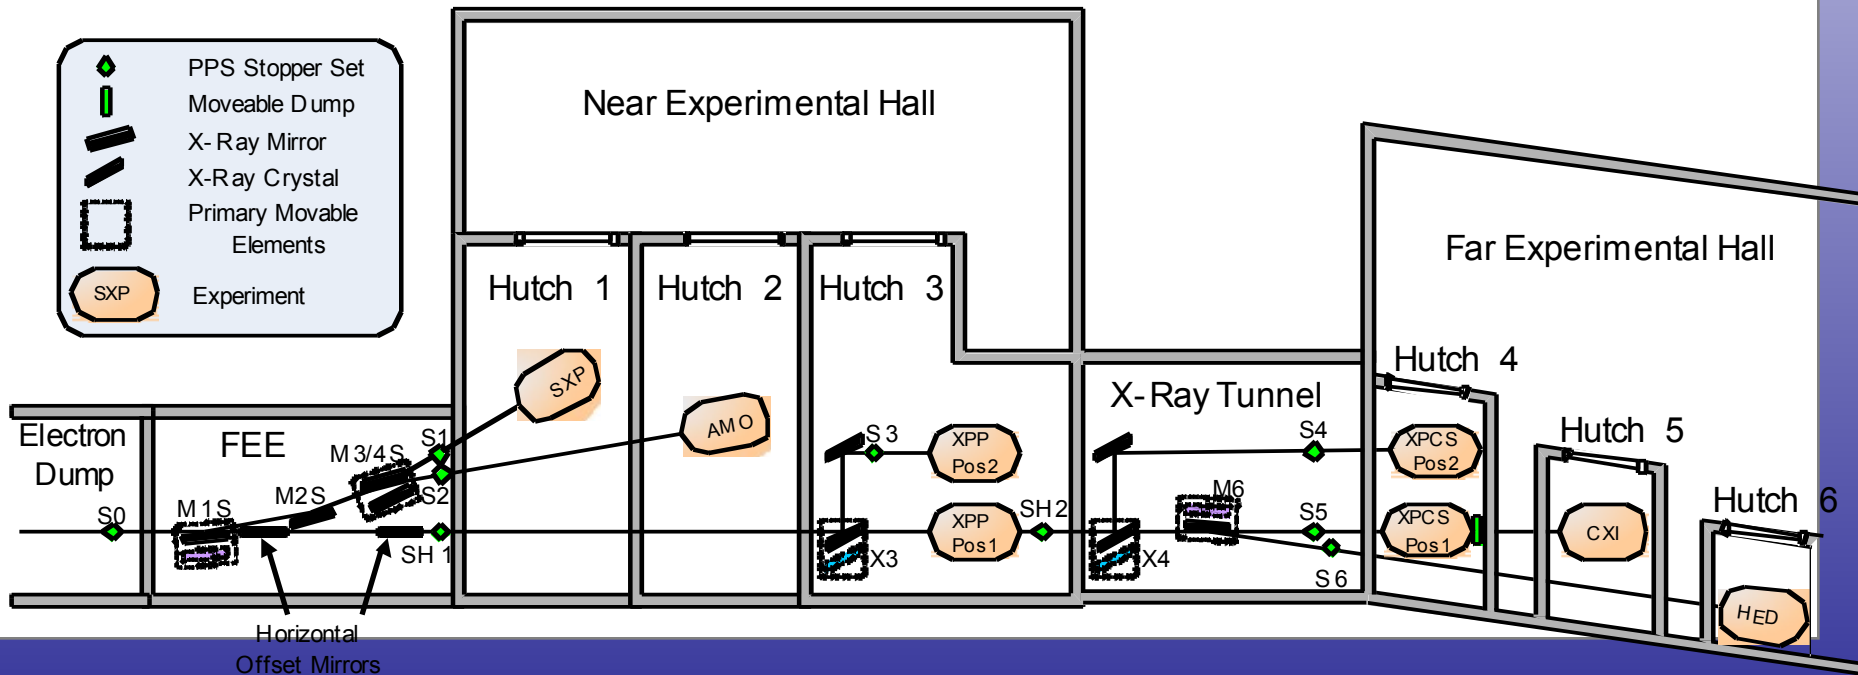
LCLS Online and Offline Computing

Alf Wachsmann
SLAC National Accelerator Laboratory

alfw@slac.stanford.edu



- Linac Coherent Light Source (LCLS) is being built right now
- It had its first LASER light on April 10, 2009
- First users on October 1, 2009



- 10 experiments at AMO
- Several involved the same people
- Many had participation from SLAC people
- 2 different experiment chambers
 - Simple one built at SLAC
 - CAMP built at ASG (Munich) and CFEL (DESY)

instrument	experiment	status	begin_time	end_time	registration_time	description
AMO	amo01109	completed	2009-12-10 08:00:00	2009-12-17 18:00:00	2009-10-09 12:38:15	Coherent diffractive imaging....
AMO	amo02809	completed	2009-12-03 08:00:00	2009-12-08 08:00:00	2009-10-19 12:11:00	Ultrafast imaging of X-ray excited clusters....
AMO	amo00409	completed	2009-11-24 08:00:00	2009-12-01 08:00:00	2009-10-09 12:31:19	CFEL-ASG Multi-Purpose (CAMP) End Station Commissioning...
AMO	amo02609	completed	2009-11-13 08:00:00	2009-11-18 08:00:00	2009-10-19 12:04:36	Explosions of clusters in intense X-ray pulses....
AMO	amo00209	completed	2009-11-05 08:00:00	2009-11-10 08:00:00	2009-10-09 10:13:09	IR laser and FEL x-ray pulse cross correlation using electron energy sid...
AMO	amo01709	completed	2009-10-29 08:00:00	2009-11-03 08:00:00	2009-09-30 23:24:42	Resonant nonlinear X-ray processes at high X-ray intensity....
AMO	amo02709	completed	2009-10-22 08:00:00	2009-10-27 08:00:00	2009-10-07 17:59:01	X-ray multiple ionization of impulsively aligned molecules....
AMO	amodaq09	completed	2009-10-22 08:00:00	2010-01-01 00:00:00	2009-10-22 11:56:33	DAQ software testing...
AMO	amo01509	completed	2009-10-15 08:00:00	2009-10-20 08:00:00	2009-10-06 20:21:12	X-Ray Non-Linear Physics Studies of Molecules with Intense Ultrafast LCL...
AMO	amo02109	completed	2009-10-08 08:00:00	2009-10-13 08:00:00	2009-10-06 19:51:06	Strong-field multiphoton processes in the high-frequency limit....
AMO	amo01809	completed	2009-10-01 08:00:00	2009-10-06 08:00:00	2009-09-30 18:26:25	Tracking transient atomic states produced by ultraintense x-ray pulses...
AMO	Commissioning	completed	2009-08-17 20:13:37	2009-09-30 20:13:37	2009-08-17 20:15:00	AMO commissioning phase...
AMO	Installation	completed	2009-08-13 15:00:00	2009-08-17 14:59:18	2009-08-13 15:54:51	AMO Installation...

- LCLS was operating at 10 Hz (mostly) or 30 Hz
- AMO produced a modest amount of data
 - Approx. 100 TB in 3 months
- Even more sophisticated users did **not** take their data home
- Improvised analysis computing on spare ATCA blades (from Level 3 DAQ)
- LCLS needs to engage with Computing Division

- Dec 19, 2009 to Apr 2, 2010: Down time work.
 - Apr 3-4, 2010:
RF process Gun, L0A, L0B. Beam to LCLS Injector.
 - Apr 5-8, 2010:
Beam to Linac and BSY.
 - Apr 9 to May 5, 2010:
Beam to LTU, Undulator and FEE.
 - Apr 17, 2010:
Beam to FEH for CD-4 milestone.
 - Apr 30 to May 5, 2010:
Five shifts of beam to AMO and SXR.
 - May 6, 2010:
AMO user run begins. SXR commissioning starts.

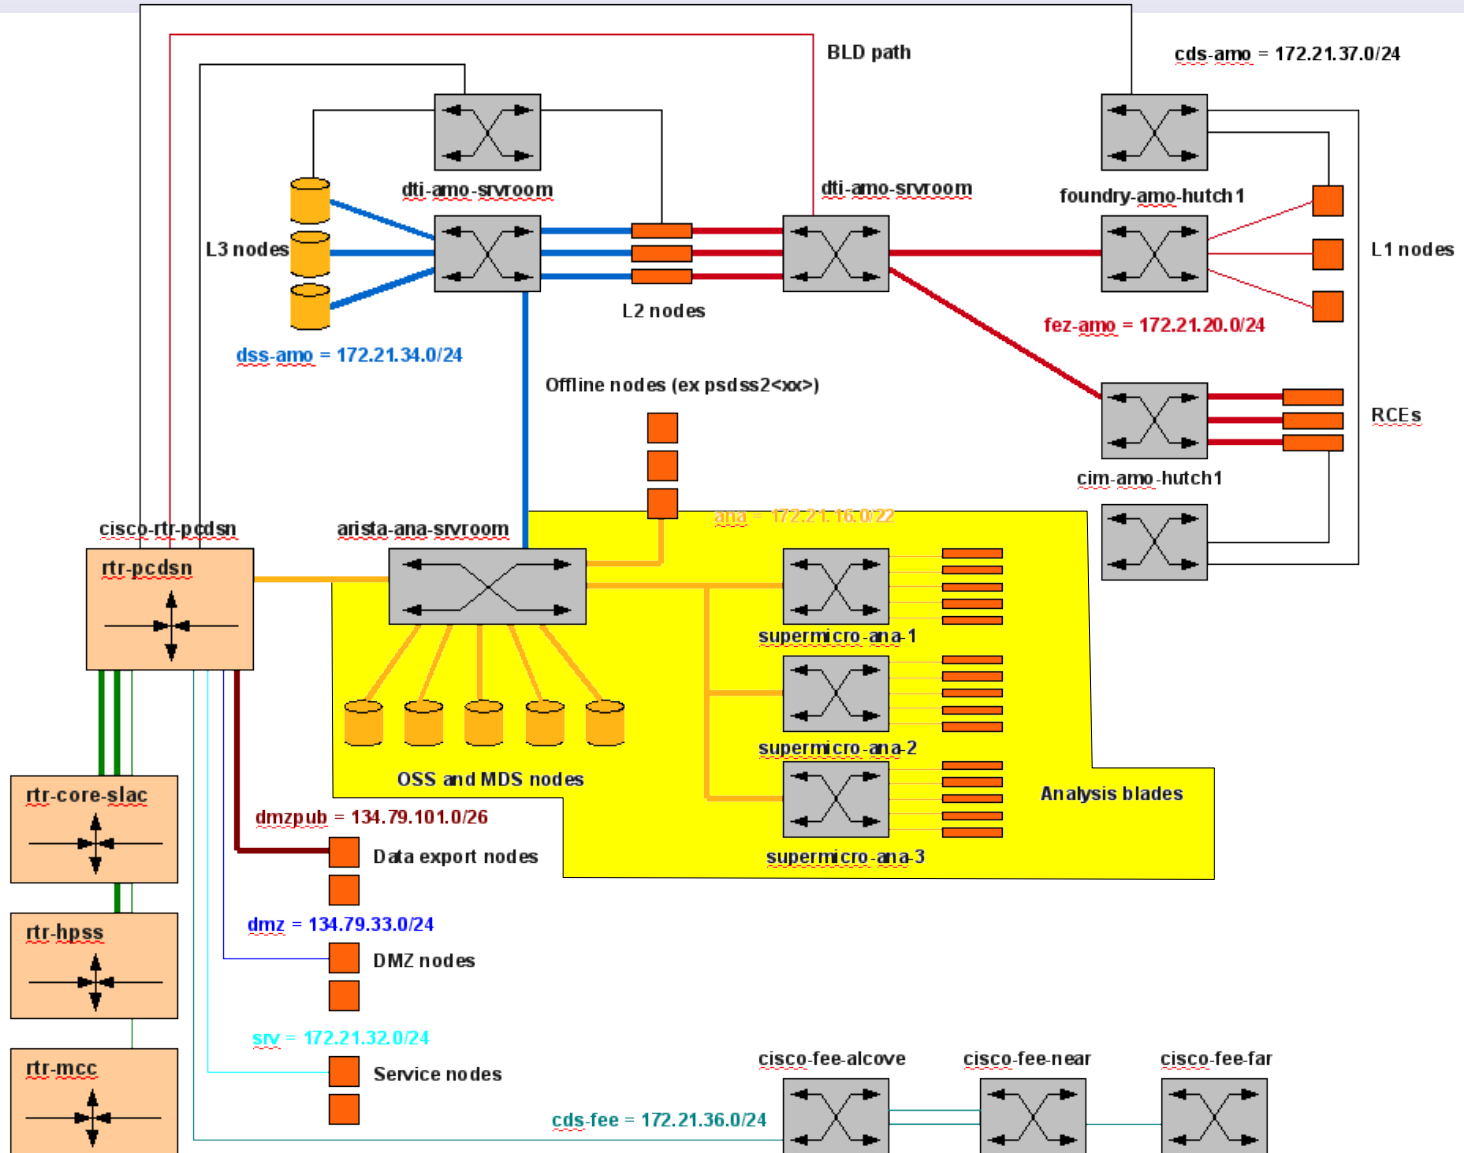
- AMO will produce more data
 - CAMP's image devices
- “User assisted commissioning” of SXR starting in May
 - About the same data amount as AMO
- Early XPP commissioning
- LCLS will operate at increased repetition rate:
 - 30 Hz, 60 Hz, and, maybe, even the full 120 Hz

- Online side had already separated “online” from “offline” Lustre
 - Users were using part of online's computing for offline purposes
- Online will have separate storage for each instrument
- Online got rid of Lustre
 - Performance problems over 10GigE network (Broadcom NetXtreme BCM5715) on ATCA blades
 - Simple rsync scheme

- LCLS upper management has decided to
 - provide all the necessary storage for data
 - Retention: 1 year
 - provide long-term archival storage
 - Expand existing SL8500s
 - Retention: 10-15 years
 - provide all the necessary offline/analysis compute power
- at no cost for users

- Data volume in 2010: 1–4 PB depending on LCLS' repetition rate
 - Will be in Lustre
 - Estimated 5GB/s bandwidth per 1 PB of storage
 - All Lustre storage on 10GigE; clients on 1GigE
- 2 Datadirect Network's S2A9900 Couplets
 - 2 controllers each
 - 8 x DDR Infiniband internally each
 - 600 2TB SATA disks each
- 4 OSSes and 2 MDSes: Dell R610s
- MDT: Dell PowerVault MD300

- 3 Supermicro Blade systems
 - 7RU with 10 Twin-Blades each
 - 500GB SATA disk per system
 - 480 E5520 cores with 3GB of memory/core
- Arista layer 2 10GigE switch
- Former conference room in NEH was retrofit with 8 Rittal water-cooled racks (~22KW/rack)
- We expect ~100MB/s sustained write into HPSS
 - Dual-copy for XTC files
 - Single-copy for HDF5 files
 - New tape drives; maybe new movers as well



- Monitor HPSS usage (ingest rate; capacity)
- Will have to learn from early 2010 experience what to do for the other instruments (SXR, XPP, etc.)
 - Idea: replicate what we did
- Analysis software framework still under development
 - Need to learn from users what they want and how they want to use it
- Look at GPUs for image analysis