Scientific Computing Implementation

Requirements

Capacity

Scheduled Upgrades / Likely purchases

Operations

- Users - Queues, Allocations, Fair share, Trouble tickets

Workflow Tools

- Auger, Jasmine, SWIF

Management

Software Support

Requirements

For full CEBAF operations, IT and ENP provide balanced, robust, and reliable software and computational resources provisioned to support JLab publication targets.

- Computing – Auger, SWIF; PBS/Torque/Maui → Slurm
  - Distributed Architecture – both onsite and offsite resources
  - Containers - Singularity
- Storage
  - Mass Storage System – LTO, JASMine
  - Disk- Lustre, ZFS, NFS
- Networking
  - 10-40G IB, 40gE DAQ GW, 10gE ESNet

ENP Computing Coordinator provides requirements and budget to IT Scientific Computing for planning and implementation.

Current ENP Capacity

- Compute - 80M core hours per year (Broadwell)
  - 5000 cores, mix of Skylake, Broadwell, Haswell
  - also still running Ivy and Sandy Bridge (DNR)

- Disk – 1 PB ENP, of >2 PB total Lustre, ZFS, 14 GB/s
  mix of data from tape, scratch, work

- Tape – up to 200 PB library capacity, if all LTO-8
  - but LTO migration expected before that happens
  - 5 GB/s across LTO-5 / 6 / 8 generations today

- Network
  - 40gE DAQ Halls to tape stage disks; IB (FDR, QDR mix) to data movers
  - 10gE offsite to ESNet

Data Analysis Cluster – "The Farm"

~250 CentOS 7 Xeon nodes, ~5000 Broadwell cores

2018: 88 dual 20 core 2.4 GHz Skylakes, 96 GB RAM, 480GB SSD, FDR IB

2016: 44 dual 18 core 2.3 GHz E5-2697V4 Broadwells, 64 GB RAM, 1 TB HDD, FDR IB

2014: 104 dual 12 core 2.3 GHz Haswells, 32 GB RAM, dual 1 TB HDDs, QDR IB

2013: 24 dual 8 core 2.6 GHz (Ivy Bridge), 32 GB RAM, dual 1 TB HDDs, QDR IB

2012: 32 dual 8 core 2.0 GHz (Sandy Bridge), 32 GB RAM, dual 500 GB HDDs, DDR IB

- Continue to encourage multithreaded codes
- No compute node procurements in 2015, 2017 during 12GeV CEBAF upgrade
- Run 5+ year old Sandy Bridge and Ivy Bridge nodes from farm, and HPC EOL, **as DNR** …

→ **Slurm upgrade, from Maui/ Torque PBS**

- Merge all Jlab (non-USQCD) resources to keep the valleys full
- USQCD cluster already switched to Slurm this summer; Farm scheduled for Q2FY19
- Testbed is ready for full testing ahead of production usage
- Interactive debug access to compute nodes, better resource management

Data Analysis FY18

Scicomp Farm Cluster Usage (org to project view)

Change time period: 10/01/2017 - 09/30/2018

Usage (org-project) | Usage (org-type) | Usage (type-org)

All Jobs

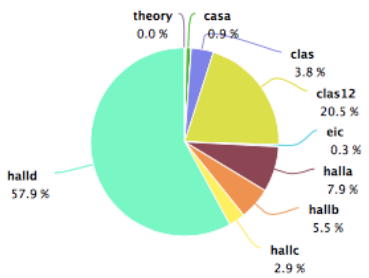

Org	Project	Job Count	Process Hour
accelerator	all	20,021	159,696
casa	all	388,709	541,598
cc	all	50	1,442
clas	all	2,063,693	2,373,554
clas12	all	874,397	12,727,505
eic	all	52,313	200,620
eshq	all	2,120	18,080
halla	all	4,129,784	4,874,866
hallb	all	2,296,699	3,427,735
hallc	all	563,811	1,774,258
halld	all	5,411,561	35,947,009
theory	all	6	17
		15,803,164	**62,046,380**

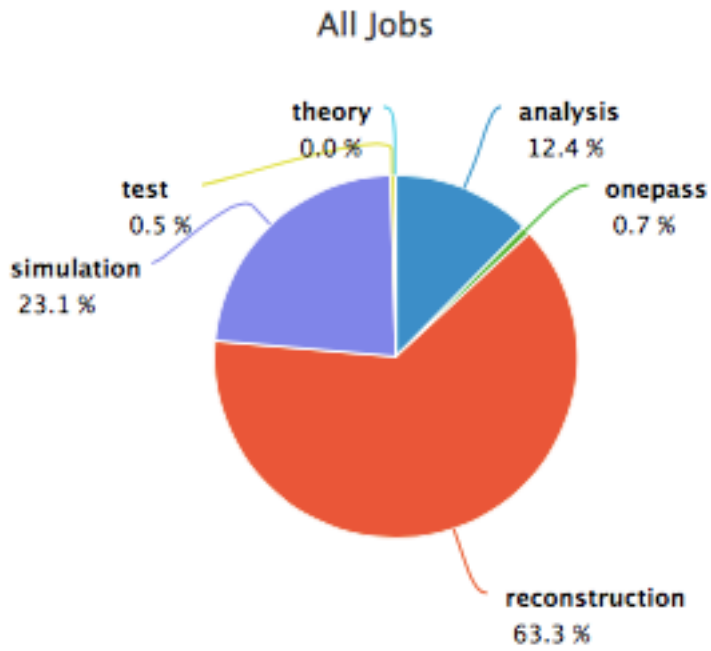

10/01/2017 - 09/30/2018 | all

Completed Job History (org)

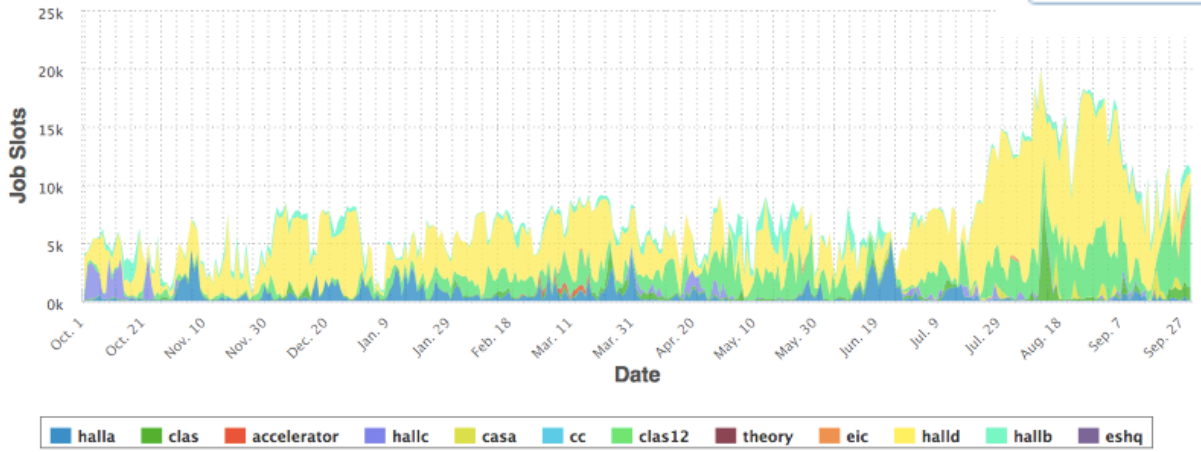

Mass Storage System – MSS

IBM TS3500 Tape Library: 12,500 slots, 24 LTO drives, 30 PB

Modular, easily expandable

11 frames: 2 library control & drives 2*340 slots, 9*1320 slot frames

8 LTO-8 (360MB/s), 4 LTO-7 + 8 LTO-6 drives (160MB/s), 4 LTO-5 (140MB/s)

16 frames max, so up to 5 more frames for more capacity and/or drives

200PB capacity at 16 frames of LTO-8

- Never have had to eject oldest media to tape vault, but could if necessary

LTO Media

Production tapes are LTO-8 (12TB/tape, 360MB/s) since Spring 2018

First batch of LTO-M8 media are on the way for testing (9TB/tape,300MB/s)

- Most cost effective LTO; avoids current Fuji/Sony lawsuit

Migration to LTO-8 – updates data from old media; frees slots

- LTO-5 underway, blank tape for reuse, then duplicate raw beam data and eject to tape vault (fireproof room in same building)

Mass Storage Accumulation

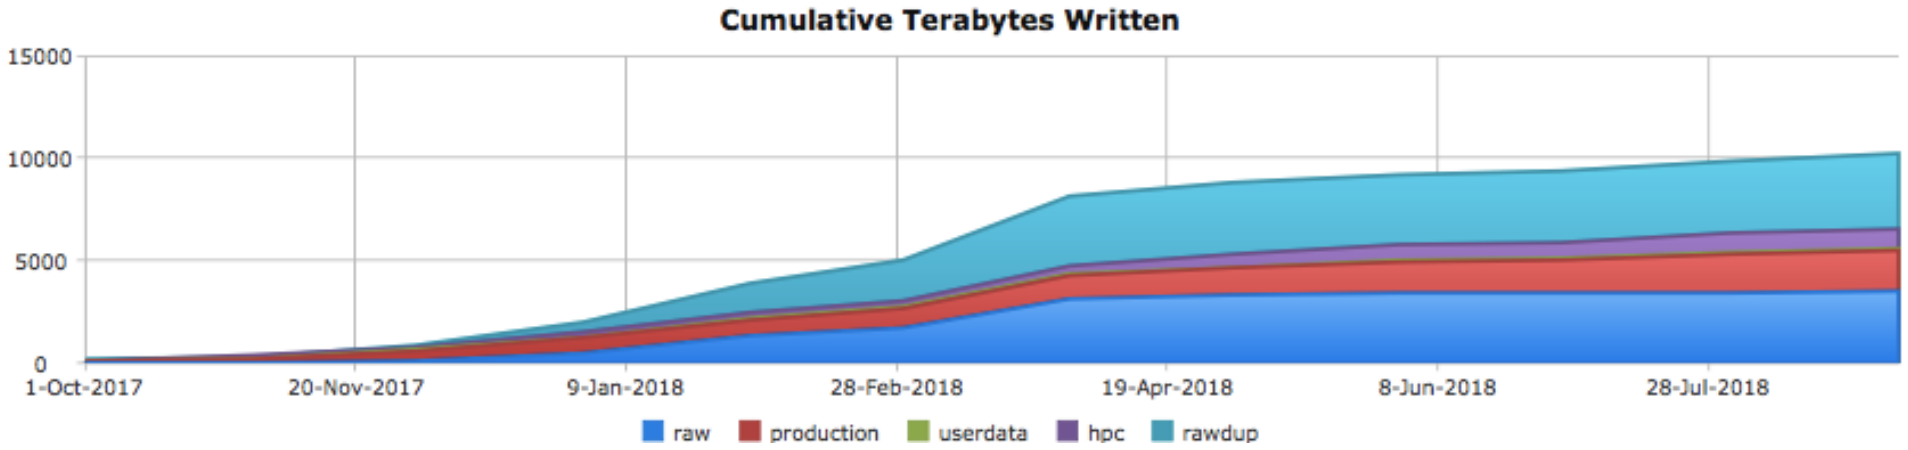

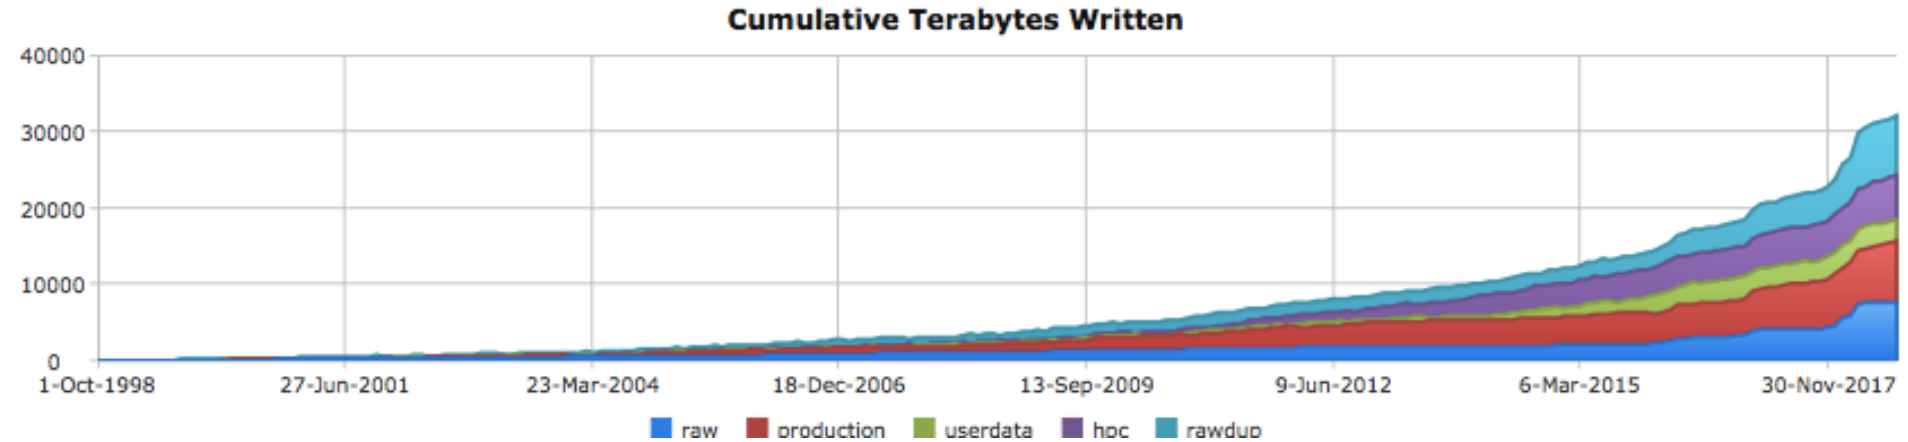

Note: duplicates of the raw experimental physics data, rawdup, are stored in the tape vault, not in the tape library

12 GeV DAQ / Online

Data storage servers in the Counting House, with tools for DAQ systems to write raw data to mass storage

- jmirror tool – duplicates the data to a staging area in the Data Center
  - Keep disk resident copy on request
- 40gE to 40gIB between Counting House and Data Center

→ Flash Data Buffer added for Fall 2018 run

- Dedicated SSDs, 14* 3 TB, 30 TB writable; 6.5 GB/s
  - avoids raw data staging on Lustre
- RAID Z2 full data protection
- 6 older DAQ gateways provide redundant backup links; do use Lustre
- Can repurpose for other high I/O tasks when beam isn't running

→ 2nd Flash Data Buffer in FY19 - increase staging pool size, redundancy

Disk-to-Tape File System

Disk area (Lustre) serves as a cache area to the tape library.

Its file system will never fill – as it approaches full, files are automatically written to tape to make room for more.

- Production since 2016
- Code already exists for HPC; ported to Physics
- Replaced Physics read-only disk cache
- Allows users/projects to write files into the area, and the files will be migrated to the tape library automatically after a certain period and as more disk space is needed
- During the waiting period, the files can be modified or deleted
- Set of command line tools to check the status of files in the area
- Modification of a file retrieved from the tape library will require user concurrence before the modified file is put back into the tape library

Fair Share Allocations

- Hall priorities are determined by ENP
- Non-ENP projects grow and shrink, based on lab priorities
- Basis for computing and disk storage

Updated for FY19 :

- Hall A 10% shared
- Hall B 45% (30% FY18)
- Hall C 10% shared
- Hall D 45% (60% FY18)
- Other: 10% shared
  - » Theory, EIC, CASA, CFDFAC, Bubble, misc

Workflow Tools – Auger

JLab software, layer above the resource manager and scheduling system; in use during the 6 GeV era

- throttles job submissions
- gathers needed data to/from tape
- preservers User environment

- But also
  - Blocks features of scheduler/resource manager
    - Interactive debugging
    - Memory specifications

Works with SWIF, but SWIF doesn't need

Auger no longer required for Slurm

Workflow Tools – Jasmine

- JLab software, interface between MSS and disk
  - Moves data from online DAQ to MSS
  - Moves data between tape and disk for farm and users
  - Manages disk to never fill cache
  - Schedules tape resources for priorities
    - Raw data copies from DAQ
    - Duplicate raw data
    - Farm jobs
    - HPC jobs
    - Users
    - Media migration

Workflow Tool

SWIF: Scientific Workflow Indefatigable Facilitator

Version 2: User Testing mode this fall

- Suspend/resume, optionally with graduated or phased job release
- **Support offsite resources**: NERSC today; OSG, then cloud
  - NERSC runs end-to-end; working through bottlenecks

Version 1: Production since 2015

- Group many jobs into a named workflow
  - Jobs can be Auger job or defined through swif add command
  - Jobs can be easily added, deleted, canceled from a workflow
  - Jobs will be organized according to the location of files to improve tape library throughput and to reduce job wait time.
- Command line interface monitoring
  - status of all jobs in aggregate, data files produced, unresolved problems
- Job problem resolution - Modify & resubmit failed jobs
- Save partial finished job results

→ Singularity containers provide portability

User and System Support

- **IT User Services HelpDesk** is front line during business hours
  - phone, email, trouble tickets:

    PHYSICS SOFTWARE– for the ENP software committee

    SCIENTIFIC COMPUTING – general questions and issues

    MASS STORAGE – next business day support

    EXPERIMENT SUPPORT– online support

  after hours support is informal, not guaranteed
- **IT OnCall** for after hours emergencies
  - Phone
  - shared by most IT CS staff, not necessarily a SciComp expert
- **Documentation**
  - http://scicomp.jlab.org/docs
  - http://data.jlab.org/drupal

Management

- Physics Computing Committee
  - Computing Coordinators for each Hall
  - Computing Coordinator for ENP Interface to IT
  - regular quarterly meetings
- JLab IT Steering Committee
  - Scientific Computing, Central Computing, MIS
  - regular annual meetings
- User Group Board of Directors Computer member
- MSS Data Management Plans
  - https://scicomp.jlab.org/DataManagementPlan.pdf
- Budget – mostly in ENP, Staffing – in IT
  - IT is Service Provider, with Annual Work Plan
  - 4 FTE (a mix of 4 Operations and 2 Developers)
- Regular computing reviews; last November 2016

Software Support

Physics division experts build, document, and support Physics software for the data analysis cluster

- Common Environment
- ROOT
- GEANT4
- CERNlib
- CLHEP
- EVIO
- CCDB
- GEMC
- QT
- JANA
- XERCESC

Documentation at http://data.jlab.org

Support via CCPR Helpdesk tickets

Facilities Infrastructure

Work completed in 2017 to meet DoE requirement of 1.4 PUE; doubled cooling and power capacities

- relocated everything within the Data Center; some things twice
- new core area for infrastructure services (file servers, database and web servers, and others) with generator backed redundant power

Space:

- 50% of 7400 square feet used; no space constraints

Power:

- Upgraded to support ~30 KW / rack, 225 KW in 8 racks
- Up to 68 nodes / rack

Cooling:

- Upgraded chiller pipe, from 8” to 12”

Planned Near-term Upgrades

- Farm – additional compute nodes
- Storage
  - MSS - more data movers and LTO-8 drives, frames
  - HDD and SSD disk servers; Lustre, DAQ
- **OSG** redundant submit host

→ **Slurm**

→ **Lustre 2.5 → 2.10 (or 2.12?) Upgrade**

- New MDS hardware in early 2019
  - redundant meta data servers, shared SSD disk shelf
  - ZFS, from ldiskfs for MDS; snapshot backups
- New FY19 disk servers
- Populate with data migrated from current Lustre, or tape
- In production by summer 2019

→ **SWIF v2**

FY19 and beyond…

- Add disk and cores to meet near-online / high data bandwidth offline computing needs
- Implement a strategy to offload simulation during JLab peak periods, and assist CLAS-12 in adopting (GlueX can use OSG in the near term)
- Transition to LTO-M8 for cost gains
  - $7/TB vs LTO-8 $13/TB
- Buy 2nd tape library if/when migration/compression is insufficient, or is less cost effective
- Collaborate with halls in exploring alternative technologies that look promising from a cost/event perspective

Summary

IT Scientific Computing projects increase resource efficiency and utilization

- New Slurm environment: better resource management and scheduling, scalable
- Lustre upgrade: support for growing file system demands
- Workflow tools
  - continue evolving to simplify Physicist's job processing and data management, including support for offsite resources
  - Write-through to tape file system keeps disk from filling

IT Scientific Computing is fully supporting 12GeV era computing

- Systems growth is requirements driven by the Halls. Scientific computing projections are updated by Physics to meet changing schedules and anticipated weeks of operations
- Facilities upgrades to IT Data Center support 12GeV era computing were completed in 2017, and meet DOE's energy footprint requirement

Resources can grow modularly – cores; disk; tape slots, drives, movers; network - to meet need proven staff have demonstrated skills for 12 GeV computing

Implementing alternatives for higher peak demands than can be achieved by local resources, including sharing with LQCD – NERSC; OSG; Cloud investigation