Computing Resources

Institutional Cluster and Computing Inventory

Stu Fuess / Scientific Computing Division

2019 ICAC

15 October 2019

Outline

- Institutional Cluster Model
  - Vision
  - Operation
  - Evolution
  - Financials
- Review of Resources
  - HTC, HPC, Disk, Tape
  - Will *not* be updating resource requests, as covered in prior ICAC
- FY19 resource additions
- Response to prior ICAC recommendations

Institutional Cluster Model - Vision

Institutional Cluster (IC) Definition:

An IC is hardware that is designed and operated by an institution, such as Fermilab, to fulfill the diverse needs of its user community. The idea is that an IC is similar to an investment portfolio which contains all types of an individual's investments. In this case, the IC contains clusters of machines that can each serve the diverse needs of the institutions' user community.

Institutional Cluster Model - Vision

- The IC is thus the collection of all local resources
  - High Throughput Clusters (HTC)
    - Designed for the efficient execution of many loosely-coupled tasks over a long period of time
    - E.g. CMS Tier-1, CMS LPC, FermiGrid clusters
  - High Performance Clusters (HPC)
    - A collection of many separate servers (nodes), which are connected via a fast interconnect
    - E.g. LQ1 cluster, LQCD clusters (pi0, pi0g, Ds, Bc), Wilson cluster
  - Storage
    - CMS and Public (all else) enstore tape systems
    - CMS and Public dCache instances, CMS EOS, NAS instances disk systems
  - Network
    - LAN and WAN portals

Institutional Cluster Model - Operation

- Users interact with a “scientific gateway”
  - Specifying workflow/job requirements, special allocations
- Gateway (= HEPCloud) can direct work to local and/or remote resources
  - These can be shared, owned, allocated, purchased, …
- Allocations, usage, costs tracked
  - Internal users == those supported by the funds used to procure & operate facility
  - External users == those who “pay” to use the IC
  - Opportunistic users == those who may utilize otherwise idle resources

All of these must be accounted for in planning, e.g. by a CRSG

Institutional Cluster Model - Evolution

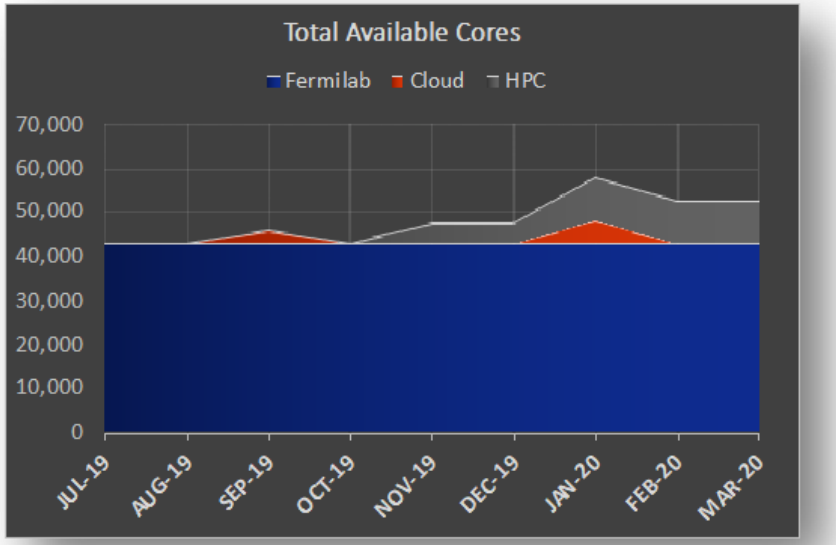

HPC Clusters

SLURM Instance #1: LQ1, lq1gwn

SLURM Instance #2: pi0, pi0g

SLURM Instance #3: Wilson

NOW

HTC Clusters: Fermi Grid, Tier-1

1. Need to have CE in front of LQ1+Wilson.
2. HEPCloud project develop Decision Engine.

HEPCloud project develop Decision Engine.

Institutional Cluster

SLURM Instance #1: LQ, Fermi Grid

SLURM Instance #2: Wilson-IC, Tier-1

Fermilab

Institutional Cluster Model - Financials

- Local resources will require a cost model
  - Should encompass all costs
    - Facility, hardware, maintenance, operations,…
  - Needs to be flexible enough to account for hardware and job type variants
    - HTC vs HPC workflows
    - Cache vs static, disk vs tape, fast vs slow I/O storage classes
    - QoS levels
  - But not so detailed as to become unwieldy
    - The “limited menu” model
  - Working on this now…
    - Where it’s hard:
      - separating “total” from “incremental” costs
      - Incorporating large “step functions”, e.g. tape technology changes

Goal: “cost recovery” from external users, but “cost awareness” from local users

Review of resources - What we have now, what is new

- The following slides show:
  - Processing resources
    - CMS Tier-1 and LPC, FermiGrid, HPC clusters
  - Processing resource utilization
    - CMS, FermiGrid, LQCD
      - Wilson cluster, by definition, is for development so utilization is sporadic
  - Storage resources
    - Disk and Tape
  - FY19 additions
    - Showing all resource types, > 95% of entire FY19 M&S expenditures

Processing: Summary of current resources

- CMS Tier-1 and LPC: to meet pledge and provide analysis platform, ~27K cores, 285 kHS06
- FermiGrid: Intensity Frontier and other HTC usage, ~19K cores, 200 kHS06
- HPC clusters: allocated, high speed interconnect (IB), some GPUs
  - No longer allocated to LQCD program:
    - pi0 : 314 nodes, 5,024 cores --- Will move ½ to FermiGrid, ½ to Wilson cluster
    - pi0G : 32 nodes, 128 K40 GPUs --- Becomes component of Wilson cluster
    - Bc : 224 nodes, 7,168 cores --- All of these are ancient
    - Ds : 168 nodes, 6,272 cores | keep for LQCD opportunistic use
    - DsG : 20 nodes, 80 Tesla M2050 GPUs --- Operate in a DNR mode
  - Newly acquired, operated as IC component:
    - LQ1 : 112 nodes (Cascade Lake) + 5 nodes with dual Voltas --- 82% LQCD allocated
- Wilson cluster: development with various accelerators, small HPC
    - Very old nodes to be retired

Processing: Last year utilization - CMS

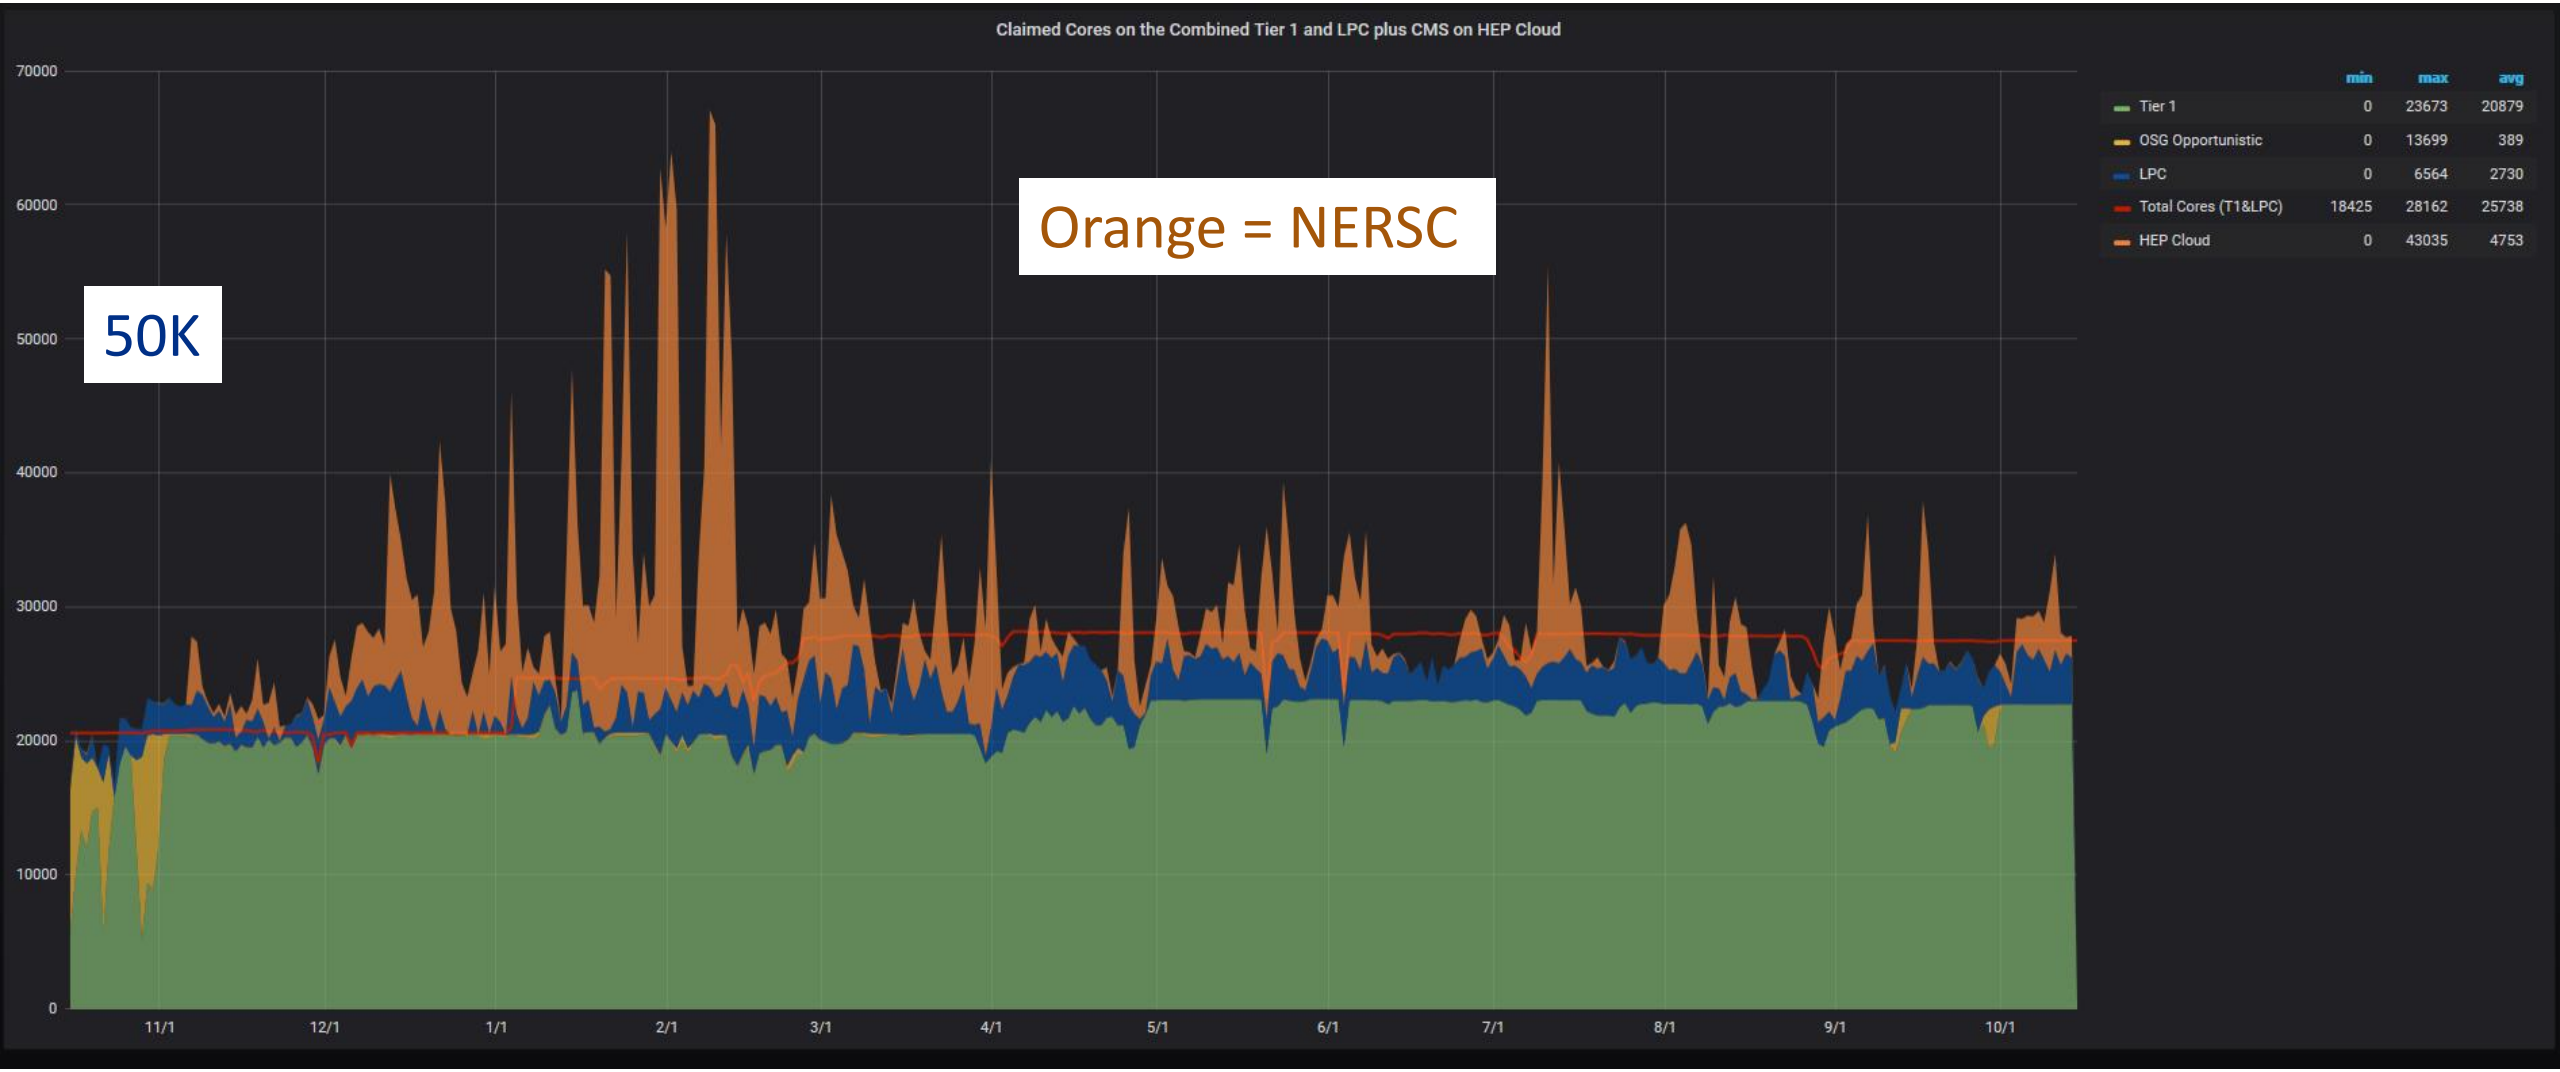

Fermilab

Processing: Last year utilization - FermiGrid

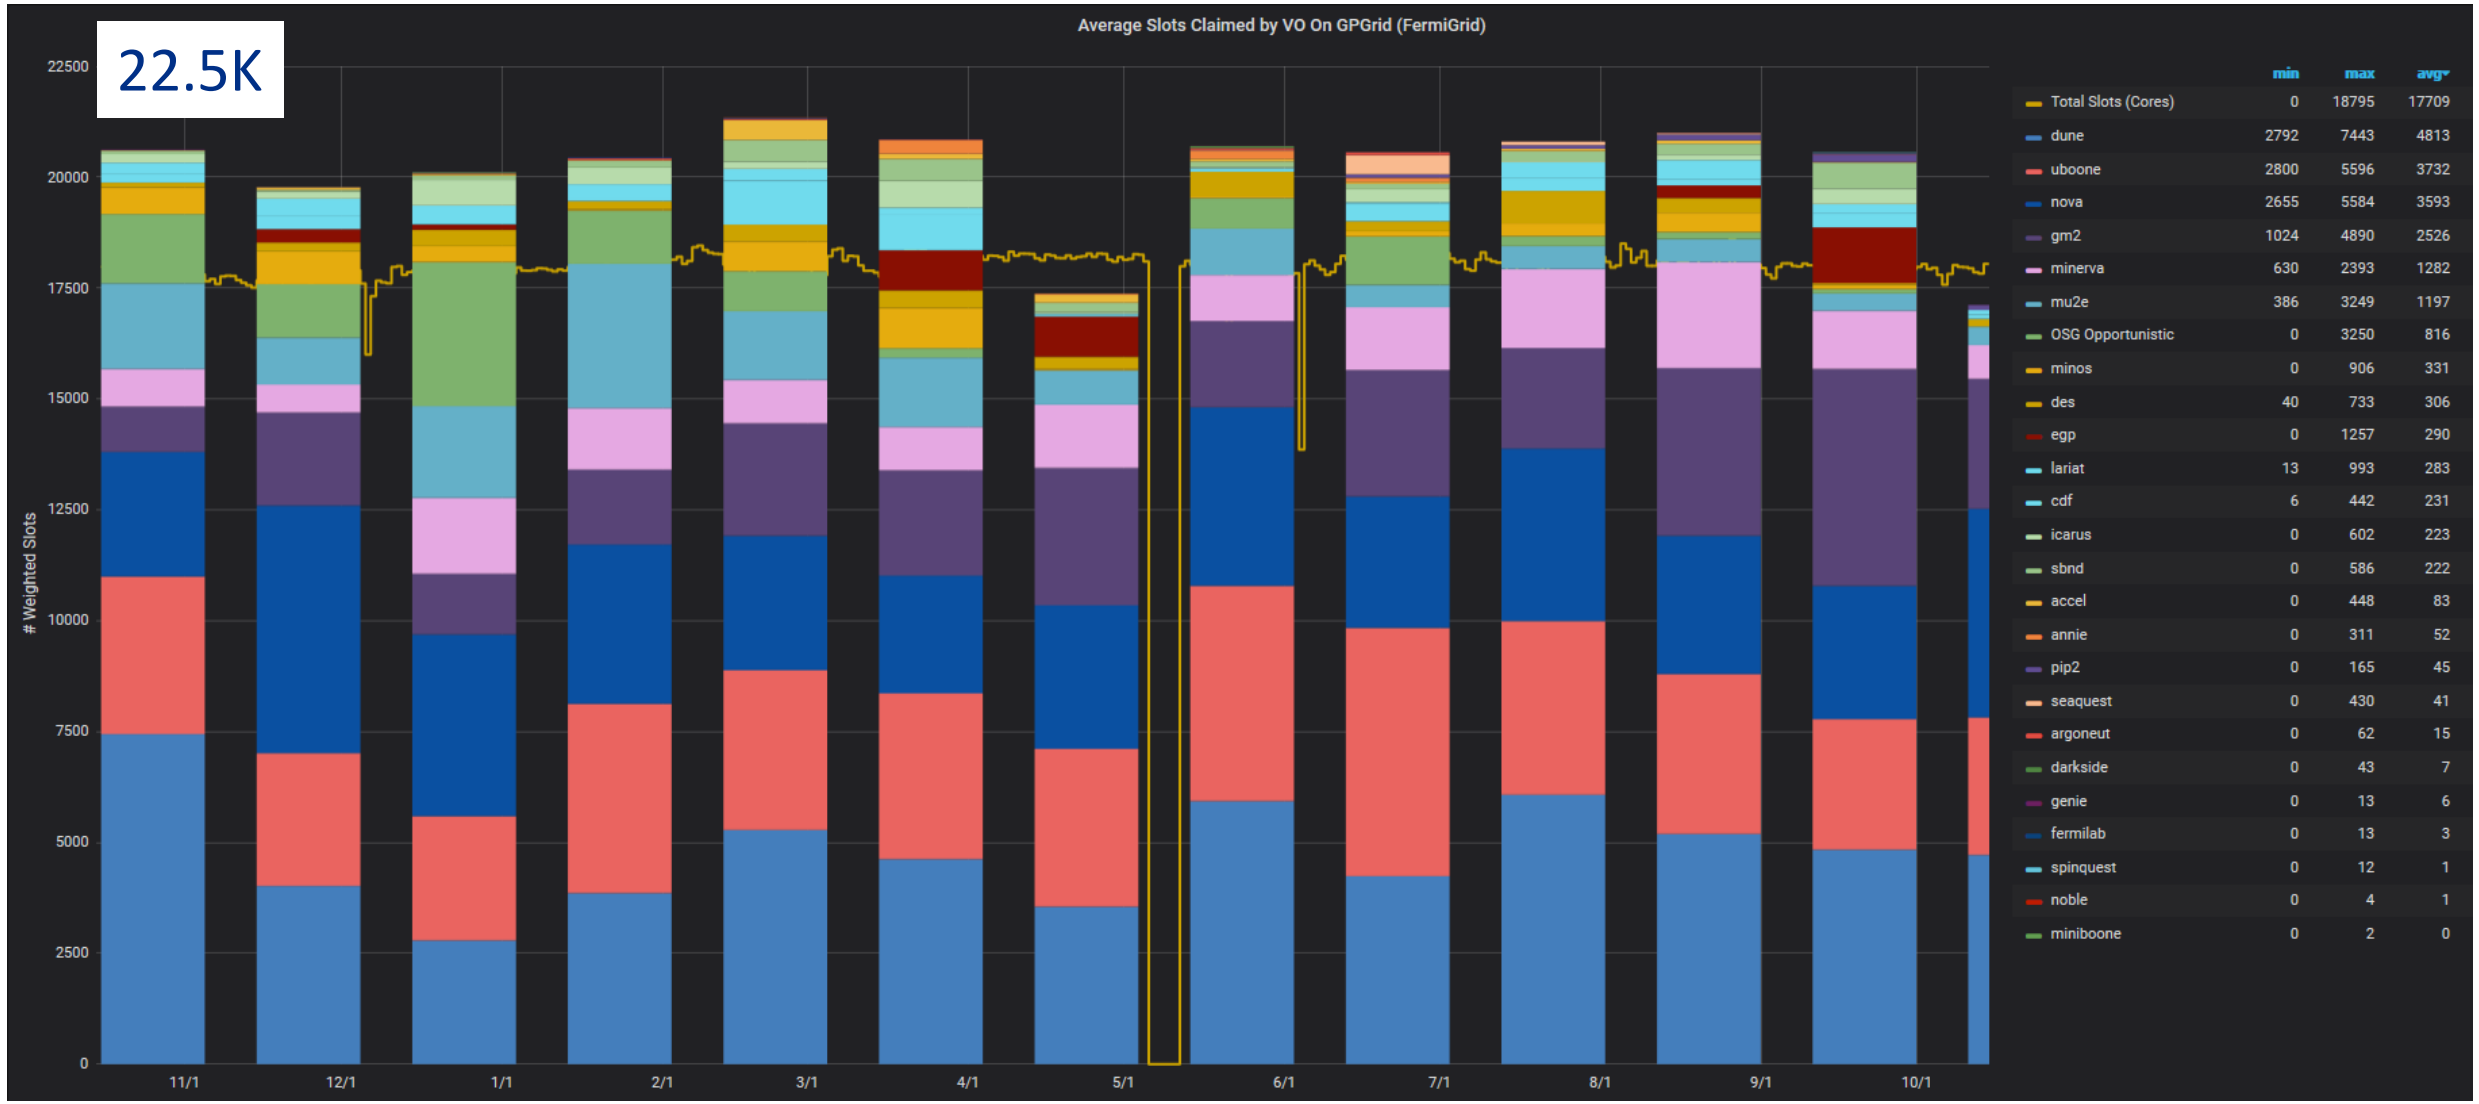

Processing: Last year utilization – HPC Clusters

pi0 314 nodes -> 105.5K node-hours/14-day bin

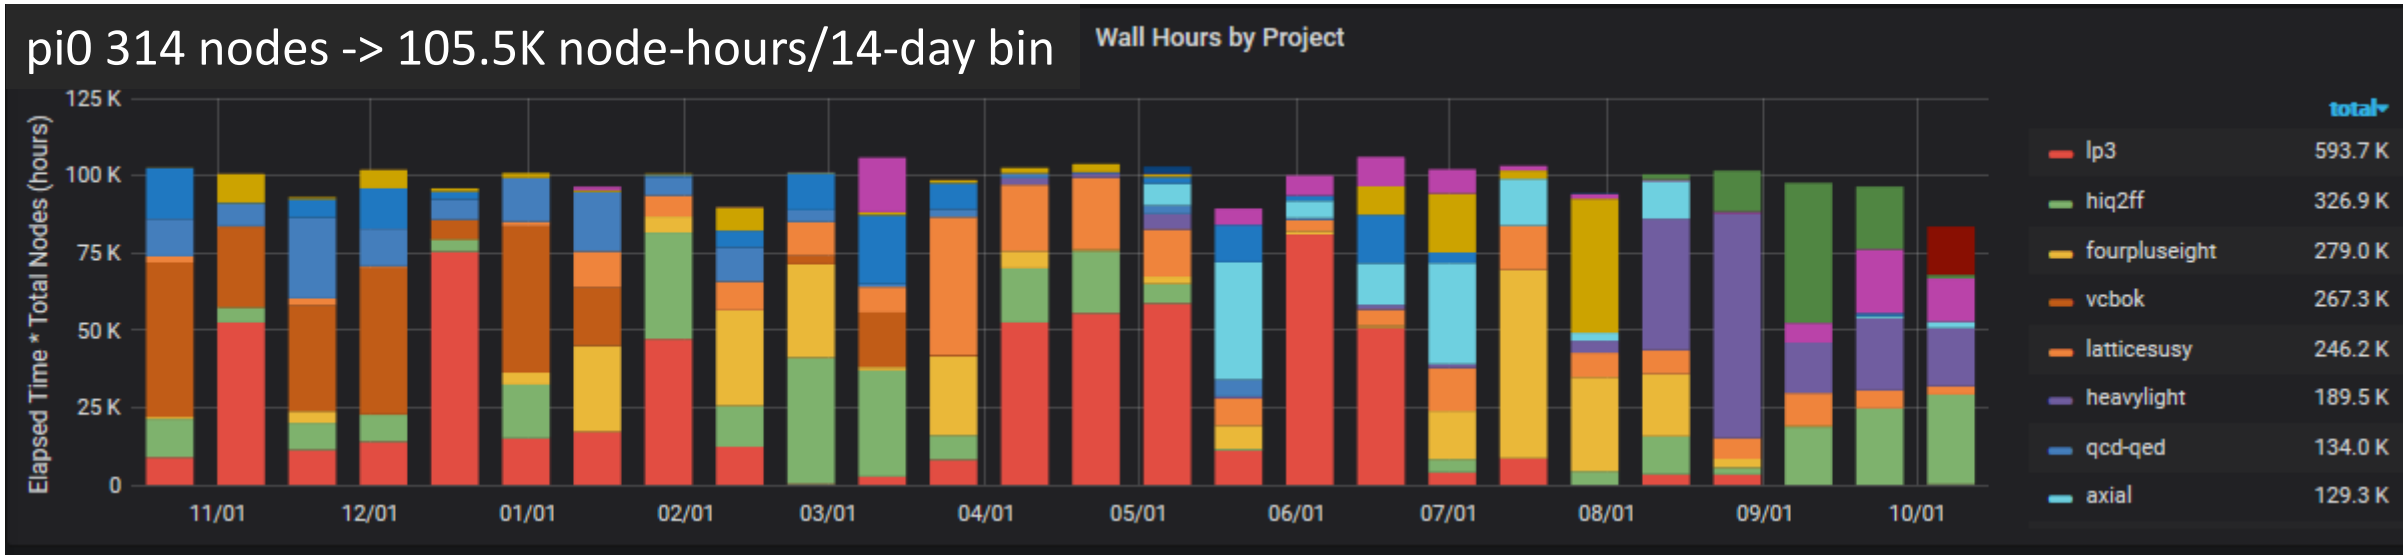

Pi0,Bc,Ds = 706 nodes -> 237K node-hours/14-day bin

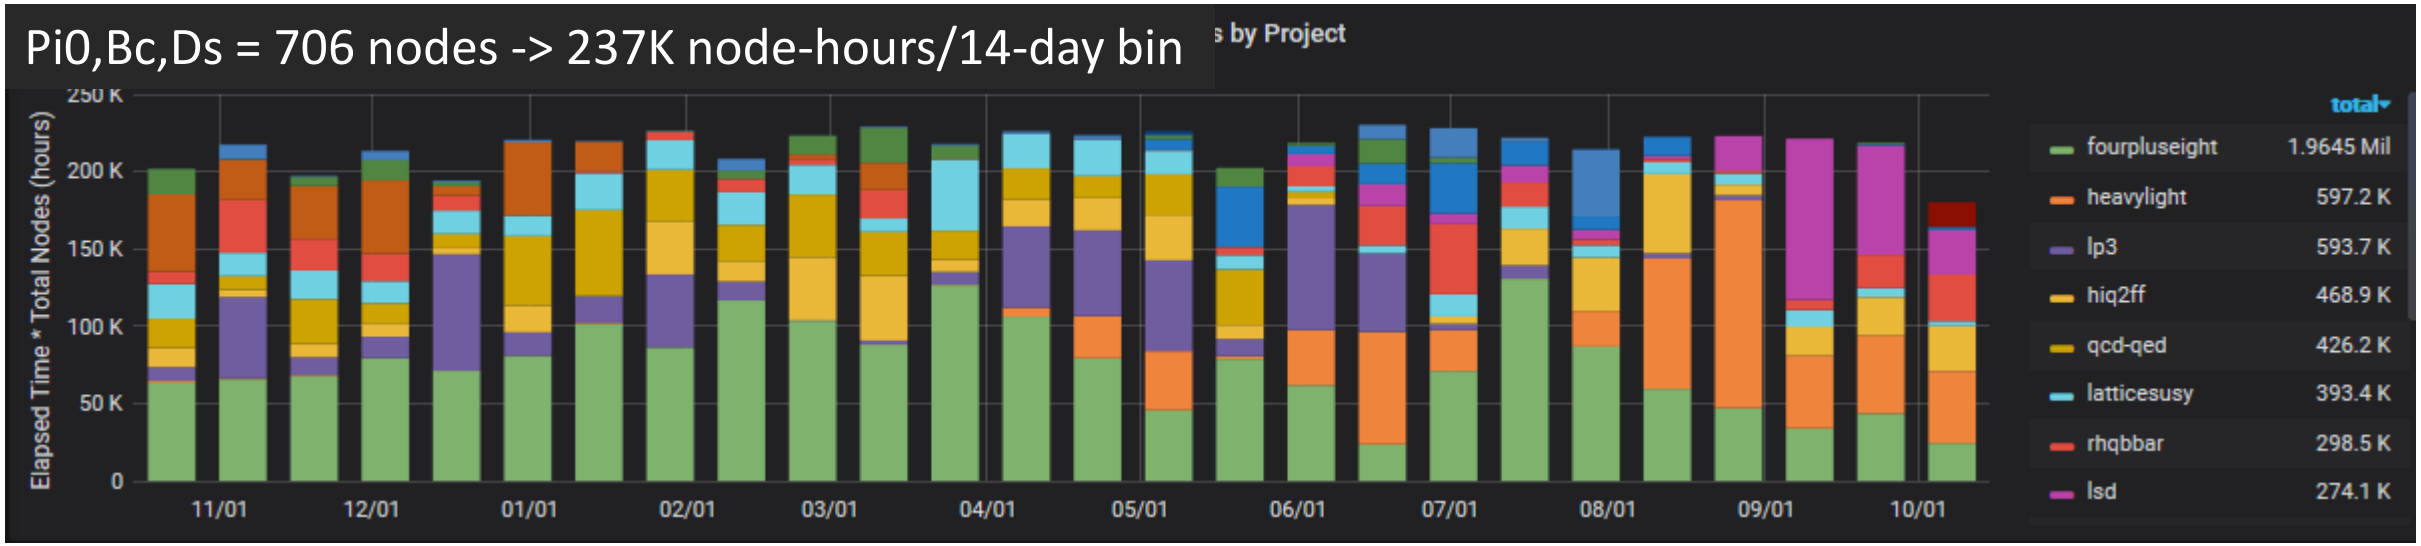

Fermilab

Storage: Current usage

- CMS

EOS 6PB

Disk dCache 27 PB

Tape dCache

Tape 68 PB
(81 PB including copies)

- Public

NAS 2PB

Dedicated dCache 4PB

Scratch dCache 2PB

Tape dCache 6PB

Tape 109 PB
(117 PB including copies)

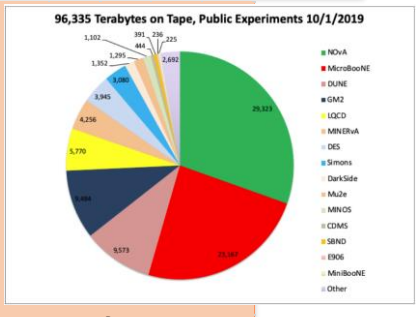

Fermilab

FY19 Resource Additions - Details

- M&S expenditures (direct costs, add ~29% overhead)

Thrust	Category	Total	Comment
CMS	Disk	$ 1,449,900	Approx 14 PB (raw)
	Network	$ 104,307	~1/2 maintenance. Adding 10 and 100 Gb capacity
	Tape	$ 474,657	Mostly maintenance. $200K media (32 PB)
Common	Service	$ 21,690	3 K8s hosts
	Tape	$ 17,200	4 Migration servers
Cosmic	Disk	$ 99,030	Approx 600 TB (raw)
	HTC	$ 86,000	4 large memory nodes
	Tape	$ 10,445	Media (1.6 PB)
HPC	HPC	$ 1,120,162	LQ1 cluster
Neutrino, Muon	Database	$ 23,000	2 Postgresql servers
	Disk	$ 366,081	Approx 3 PB (raw)
	HTC	$ 344,400	~5.4K hyperthreads in 42 Epyc Rome nodes
	Interactive	$ 39,000	Refresh VM infrastructure
	Network	$ 185,924	~1/2 maintenance. Adding 10 and 100 Gb capacity
	Service	$ 40,750	Refresh server infrastructure
	Tape	$ 1,036,545	2nd IBM TS4500 LTO8 library. $450K maintenance. Media (38 PB)
R&D	Disk	$ 373,200	18 systems, ~660 TB (raw) NVMe disks
Grand Total		**$ 5,792,291**	

FY19 Resource Additions - Highlights

- CMS
  - Adding ~14 PB of disk
  - Purchased 32 PB of media (all M8, 9 TB/cartridge)
- Neutrino/Muon (aka Detector Operations)
  - Adding ~3 PB of disk
  - Adding over 5K FermiGrid “slots” (42 Epyc Rome nodes)
  - Commissioned 2nd IBM TS4500 100PB capacity (with LTO8) library
- HPC
  - Commissioned 112 node cluster (Cascade Lake, 100 Gb Omnipath)
- R&D
  - Purchasing sizable NVMe disk arrays (multiple sizes, net 660 TB)

Response to related ICAC recommendations (my paraphrasing of such…)

- *Fence computing resource funding from, or within, the Detector Operations B&R*
  - We are still at the mercy of general DetOps budget, and especially competing demands for salary funding (as representative of needed effort)
- *At least get a manageable level of funding*
  - Approx. $900K was redirected to hardware resources from within the sector in FY19
    - Only a one-time possibility
  - Allowed refresh of worker nodes and disk R&D purchases
- *Avoid fencing of resources*
  - Institutional Cluster model addresses this, particularly local HPC resources
- *Plan for system repurposing or retirement*
  - A large number of Lattice QCD systems are no longer allocated to the project
    - Will redirect a number to FermiGrid, make others openly available on Wilson (HPC) Cluster

Conclusion

Additional details available if requested