

# Hot Issues at SLAC Analysis Facility.

Wei Yang

# New Green Field @ SLAC Computing

SLAC will build a new Scientific Data Facility (SDF)

- Decoupled with the existing computing environment
  - To free from the current constraints inherited from decades of HEP computing.
- To address the need to Lab wide computing needs
  - Open to stakeholders: LCLS, CryoEM, Fermi Gamma, Accelerator simulation, (ATLAS), etc.
  - Lab contributes a portion of CPU/GPU and disk and make them available to all users.

Technical Changed:

- LSF → SLURM
- HTC → HPC with throughput in mind
- Lustre from DDN for home (SSD), scratch data space, project spaces
- Windows AD for authentication
  - slowly remove dependence on Unix authentication - toward a single SLAC computing account

# Challenges and opportunities for SLAC T3

For a while we will have two computing environments

- Until the exist one is phased out
  - Existing hardware will likely be obsoleted in ~ 3 years
- SDF with SLURM will likely have much better support in GPU scheduling
  - Some users already started using GPUs at SDF
  - Need to build Grid interface to SDF/SLURM
- **When SDF is online, we want to put our new users on SDF side by default**
  - **They will scavenger petition, which may be inadequate** (unless ATLAS invest in SDF), So
  - We would allow existing or new users to use both if they want
- Original target: Online by April, in full production by summer (LCLS targets)
  - **Facing delays** -- SLAC curtailment of on-site operations starts at 3/17

# Other activities across the AFs

- Improving users tutorial document and out reaching
- Introduce Jupyter and Xcache as new services to our users
- Looking at defining a baseline Jupyter environment
  - Help users to smoothly migrate from one place to another
  - But it also help us:
    - We need both diversity to provide a broader coverage, and
    - Joint effort to improve our technical capability

# Backup Slides

## SDF Phase 1 Storage

Phase 1 storage is on-site and we are deploying it!

- Two DDN SFA18KX controllers
  - 72 drives per controller
  - 5 enclosures per controller, each with up to 90 drives
  - Up to 3.2 million IOPS, 90GB/s per controller
- ~250TB SSDs for metadata and home directories
- ~7PB HDDs for data storage
- A single namespace across multiple controllers and storage pools
  - Declustered RAID



## SDF Phase1 GPU Nodes

- We've deployed 11x GPU nodes funded by SLAC Machine Learning initiative
- Existing CryoEM GPU servers will also be migrated to SDF
- Supermicro SYS-4029GP-TRT2
  - 2x 12-core Intel Skylake Xeon Gold 5118 CPUs @ 2.3GHz
  - 192GB RAM
  - 10x Nvidia GeForce RTX 2080 Ti
    - 4352 CUDA cores @ 1.35GHz
  - 10GbE Base-T Ethernet
  - 1x 480GB + 3x 1.92TB SSDs



## SDF Phase 1 CPU Nodes


- We just ordered 80x Dell PowerEdge C6525 Servers
- Node Specs:
  - 2x **64-core** AMD *Rome* EPYC 7702 CPUs @ 2.0GHz
    - 2x AVX-256 SIMD units per core
    - **2 TFLOPS** per CPU
  - 512GB RAM (4GB per core)
  - Mellanox ConnectX-6 **100Gb/s** HDR100 InfiniBand Adapter
  - 10GbE Base-T Ethernet
  - 960GB SSD







# Fat-Tree IB Network Topology (2:1 oversubscription)

SLAC

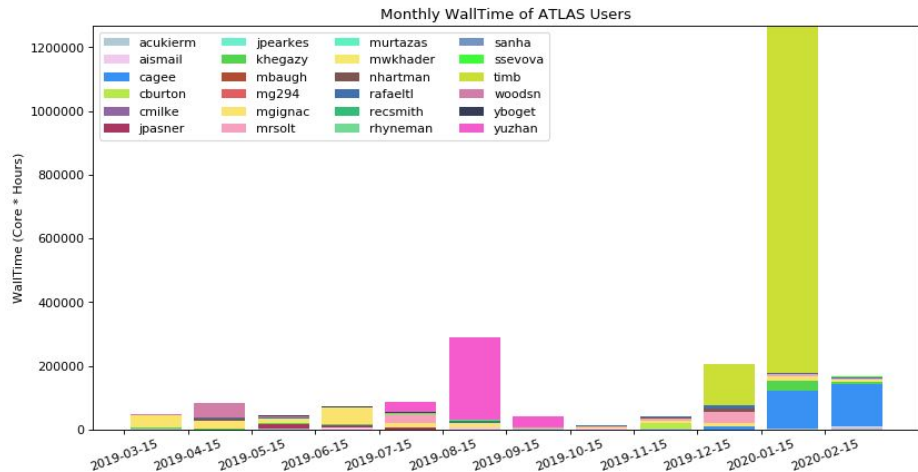
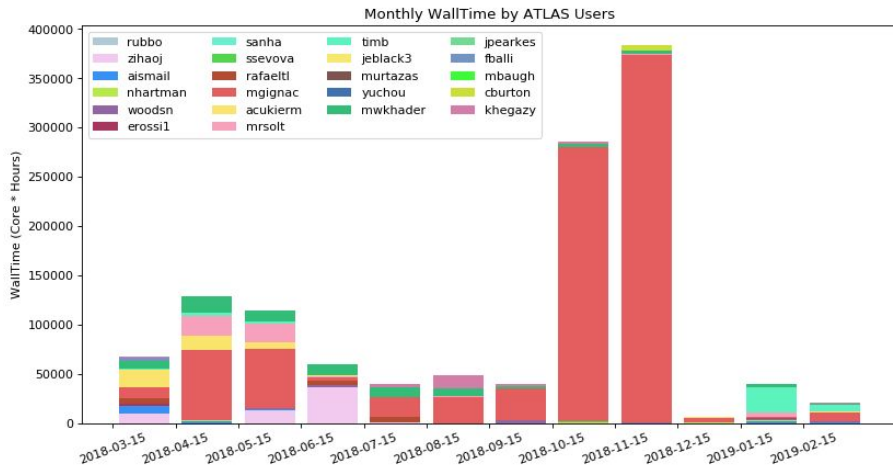


# Future Expansion

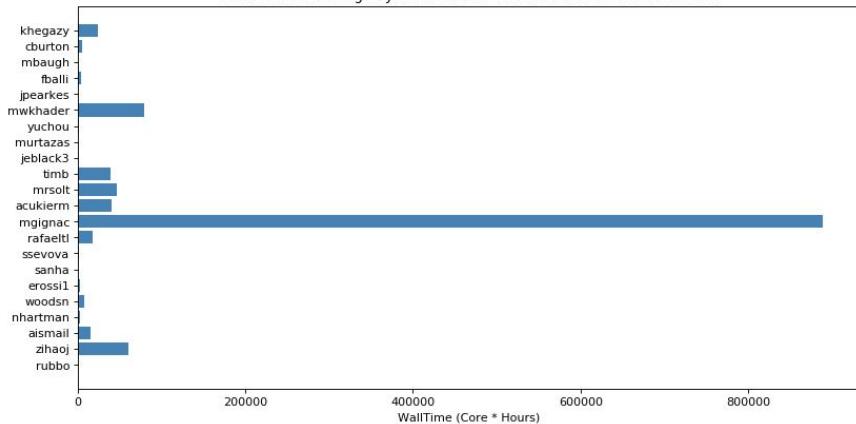
SLAC



# Batch Users in 2018 and 2019



Total WallTime Usage by ATLAS Users From 2018-03-15 To 2019-03-15



Total WallTime Usage by ATLAS Users From 2019-03-15 To 2020-03-15

