



# *Slurm User's Group 2019*

## LANL Site Report

*Joseph 'Joshi' Fullop  
Steve Senator  
Nichole Dobson*



September 17, 2019



Managed by Triad National Security, LLC for the U.S. Department of Energy's NNSA

# HPC at LANL

System	Nodes	Type
Ising	1	D-Wave
Gadget	24	Cray
Fog	32	
Kodiak	66	GPU
ViewMaster III	92	
Capulin	175	ARM
Thunder	175	ARM
Woodchuck	190	
Trinitite	200	Cray (Haswell/KNL)
Snow	368	
Badger	372	
Fire	1,104	
Ice	1,104	
Cyclone	1,104	
Grizzly	1,490	
Trinity	20,000	Cray (Haswell/KNL)

# ***Paid Requests For Enhancement***

- ***DataWarp Improvements***

- StageIn and StageOut logic
- Job cancellation considerations

- ***Resiliency Features***

- Timing of task signaling for cleanup
- Sbcast communications over High Speed (or other) Network
- Slurmdbd messages fast fail

# ***Paid Requests For Enhancement***

- ***JobComp Script additions***
  - Which Scheduler Launched Job
    - Main, Backfill or Submit
  - Job's Requested Features
  - Last 'Reason' a Job was waiting in the queue
  
- <http://bugs.schedmd.com/>
  - **Advanced Search:**
    - **Status=Open**
    - **(Detailed Bug Info) Severity = 5 - Enhancement**

# *Investigations and Solutions*

- ***File System Aware Scheduling***

- Implemented using licensing mechanism instead of GRES
- Defaults to requiring ALL networked file systems via submit filter

- ***YeTi – security zone inheritance***

- Cluster compute nodes with two independent security domains
- Late provisioning, Environment Validation and Job Requeue
- Implemented with active node features
- Prolog & Epilogue environment is now richer

# *Investigations and Solutions*

- ***Immediate (as possible) Interactive Jobs***
  - Small reservation with replaceable nodes to minimize waste
  - Utilizes Qos and Partition to limit 1 per User and Max Total Nodes
  
- ***Diminishing Account Share Fairshare Strategy***
  - Equal Access in a Usage-based Allocation World
  - Faux Bank Accounting
  - Fast Decay to adjust to dynamic demand

# *Investigations and Solutions*

- ***Orchestration of Rolling Updates***
  - Goal of eliminating significant amount of downtime
  - Dependencies on Facilities Work
  - Dependencies on System & Infrastructure Nodes & Resources
  - Dependencies of Jobs on Compute Environment
  - Management of Prior and New Workload
  - Boot and Provisioning Integration
  - Use of Active Node Features & Submit Script

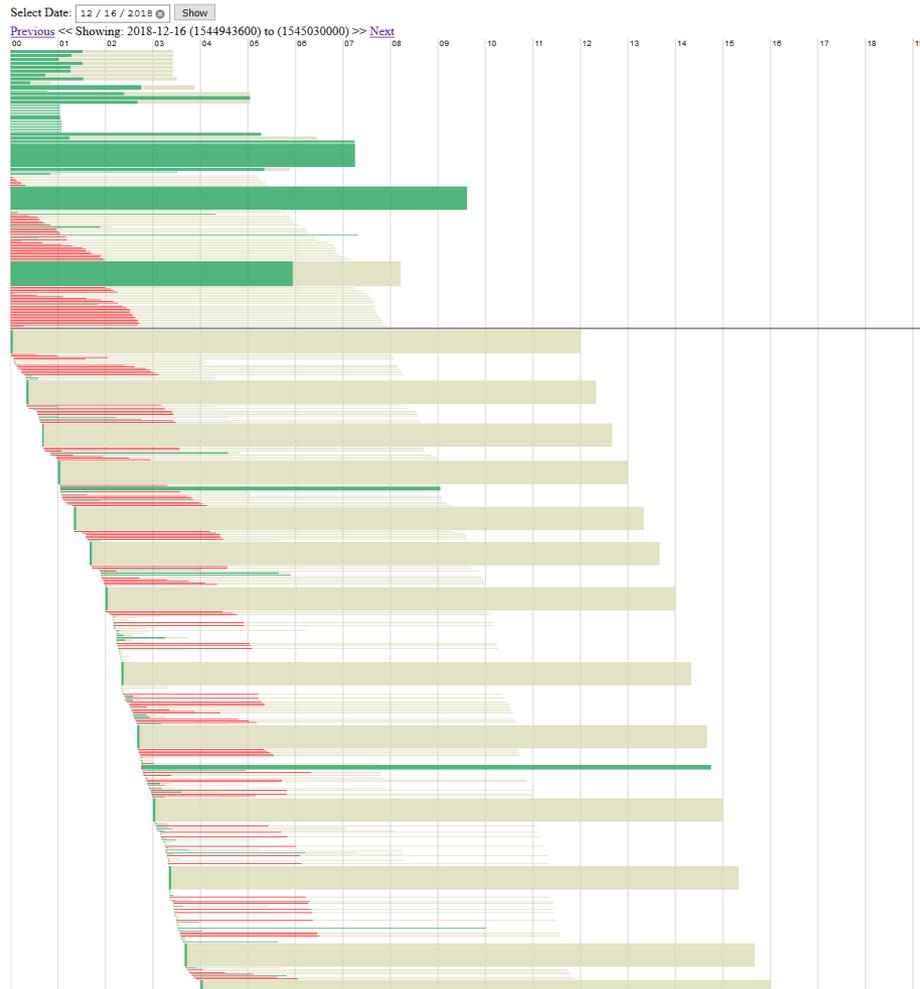
# Tools, Tech & Toys

- **Workflow Cascade Visualization**



# Tools, Tech & Toys

- **Workflow Cascade Visualization**



# ***Future Work & Collaborations***

- ***Federated/Multi-Cluster Dependencies***
- ***End-of-Job Reports to Users***
- ***Job Profiling and Deviation Detection***
- ***Task Signal Mechanism on Task Failure***
- ***Adoption of DataWarp RESTful API for Data Staging Clusters***

# LANL Workload Management Group

September 17, 2019

Slurm User's Group  
Salt Lake City, Utah

