



From Moab to Slurm: 12 HPC Systems in 2 Months

Peltz, Fullop, Jennings, Senator, Grunau

Tuesday, 26 September 2017

UNCLASSIFIED



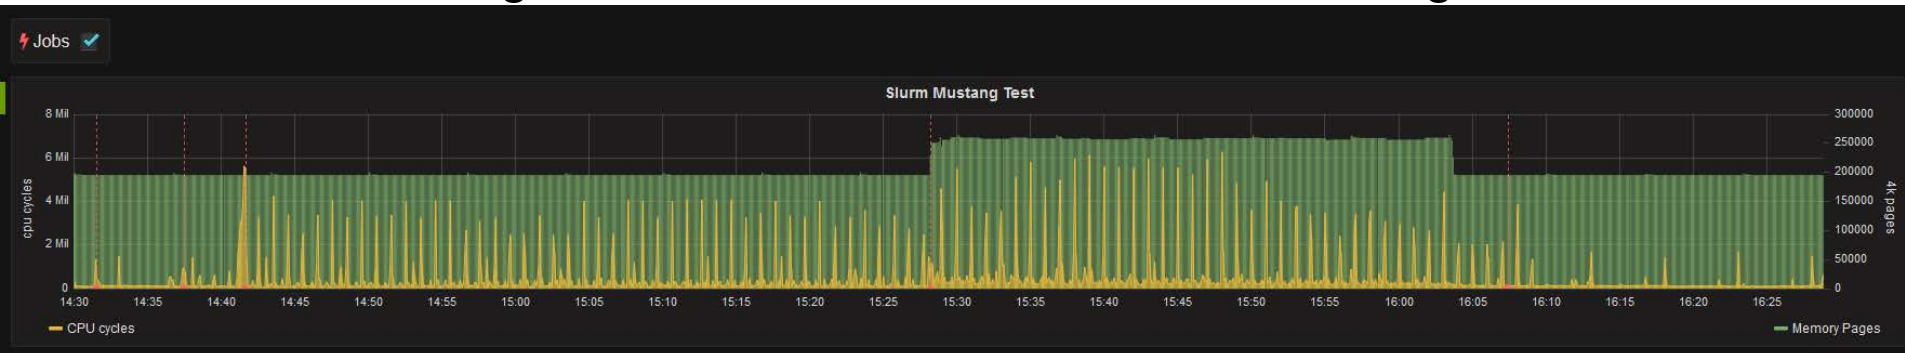
Where we started

- Multiple systems with various operating systems and architectures
- Moab as the primary scheduler
- Resource Managers included:
 - PBS
 - PBS/Alps
 - Slurm (version 2.3) – Dictated by TOSS

UNCLASSIFIED

Preparations

- Test & Development systems
 - Kit, Cub, VMs
- Scale Profiling on decommissioned cluster – Mustang




- Trip to NERSC
- Early Hybrid system - Woodchuck

UNCLASSIFIED

Systems List

System	OS	Zone	Date	Nodes
Mustang	TOSS2	Open	Jan 12	1,500
Woodchuck	TOSS3	Open	Mar 8	226
Pinto	TOSS2	Open	Apr 10	154
Trinitite	CLE	Open	Apr 21	100
Snow	TOSS3	Open	Apr 24	368
Moonlight	TOSS2	Open	May 1	308
Ice	TOSS3	Secure	May 8	1,104
Grizzly	TOSS3	Open	May 15	1,490
Wolf	TOSS3	Open	May 17	616
Lightshow	TOSS2	Open	May 22	16
Fire	TOSS3	Secure	May 24	1,104
Luna	TOSS2	Secure	May 30	1,540
Viewmaster II	TOSS2	Secure	May 31	60
Trinity	CLE	Secure	June 1	19,364

Timeline



UNCLASSIFIED

Accounting Integration

- Allocations: LDAP to Slurm
 - sacctmgr load file...
 - Additive only
 - Adds may fail and derail
 - Finding discrepancies & issuing sacctmgr commands
 - May fail silently
- Accounting: Slurm to HPCStats
 - Join slurm_jobcomp with slurm_acct_db tables to generate csv
 - Import to HPCStats

UNCLASSIFIED

Implementation

- Ansible
 - Developed playbooks
 - Removal of Moab/PBS
 - Installation of Slurm & databases
- Interstate – Tool for managing Dedicated Service Time (DST)
 - Stop Slurm services
 - Generate maintenance reservation(s)
 - Revert above for Return To Service
- Test Rollout/Rollback during regular DSTs

UNCLASSIFIED

User Education

- Self-presented training for users
- Official SchedMD training for Admins and Users
- Just In Time training scripts
 - Created and used these instead of wrapper scripts
 - Tells users to use the equivalent command instead
- Consulting Staff

UNCLASSIFIED

General Hurdles

- Firewall configs
 - Functioned, but did not scale well.
 - SYN Flood protection
- FastSchedule / Memory Specification
 - Opted to not use Memory as part of CR selection type
 - Exclusive use allocations
- Bind Affinity
 - Fun with hyper-threads and NUMA domains
- Job Accounting

UNCLASSIFIED

Architecture on Cray

- Planning for Scale
 - Slurmctld on a unused service node rather than SDB
 - Backup slurm controller/slurmdbd/mysql on external node
 - Test boots into slurm on Trinity
 - Find bottlenecks and scaling issues
 - Debug level too high
 - Slurmctld log file needs to be on local file system to slurmctld node
 - Cray's NHC was extremely slow at cleaning up jobs
 - Various timeout settings

UNCLASSIFIED

Cray Scaling Issues

- Parallel File System
- PMI
- bcast
 - Temporary solution to help synchronize launches

UNCLASSIFIED

Cray Customizations & Tuning

- MCDRAM
- TCP Timeouts
- Parallel srun
 - Specify GRES:craynetwork=0
- VM Overcommit
- DataWarp
- Re-provisioning of KNL nodes

UNCLASSIFIED

Problems Encountered Specific to Cray

- ncmd
 - TCP/IP communications break down
- aeld
- capmc
 - Slurm does not see correct node features
 - Leads to job denial after submit script supplies default features
- mapn
 - No parallel sruns

UNCLASSIFIED

Post-Migration Cleanup

- Config Consolidation
 - CR_Select type
 - Host Definitions

UNCLASSIFIED