

*Exceptional service in the national interest*



5/13/2016

# Sandia National Labs Scheduling Approach

Carol Alvin

5/17/2016



Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

# SNL Configuration

- Scheduling software is Adaptive Computing Moab
- Resource manager is SLURM
- Visualization is accommodated on the regular clusters. A set of nodes is reserved for viz.
- Interactive work is handled with a set of nodes with a short wallclock limit. Interactive jobs are not limited to these nodes.
- For most clusters, maximum run time is 96 hours. On newest cluster, 48 hours.
- All clusters are scheduled at the node level

# User Profile

- Mix of Power Users and casual or new users
  - Scientists and engineers with important deadlines
  - Code developers running regression tests
  - Students just experimenting
- Jobs are submitted via command line only

# Clusters & Resources Supported

- 10 traditional HPC clusters in 4 different security network environments
  - Several different hardware types
  - Goal is to keep OS and libraries consistent
  - Dedicated cluster for single-node jobs
- Also Open Stack data analytics clusters
- GPUS only on limited and experimental basis
- No VM/Docker services are provided
- No cloud bursting
- Multiple parallel file systems and NAS are provided
- Tape archive system

# Job Mix

- Mix of parallel and single-node
- Combination of interactive and batch, most work is done in batch mode
- No high-throughput at this time
- Queue depth varies by machine and network, generally more secure environments have fewer jobs queued

# Job Mix (cont.)

- In the last 7 months:
  - 4+ million jobs were run
  - Nearly 3 million run on single-node cluster
  - Anywhere from 8 to 30,000 processors
  - Vast majority use <512 processors
- On the two busiest clusters:

Wall Clock Limit	Avg Run Time	Avg Wait Time
48 hours	1 hours	1 hour
96 hours	1.2 hours	1.1 hours

# Capacity Computing Resources 2016

System	Vendor	Nodes	Processor cores total	Processor:Sockets:Cores/Socket	Memory/Node	Memory/Core	TFLOPS	GA Date
Sky Bridge	Cray	1848	29,568	2.6 GHz Intel Sandy Bridge:2S:8C	64	4	588	Mar-15
Red Sky	Sun	2,823	22,584	2.93 GHz Intel Nehalem:2S:4C	12	1.5	264	Apr-10
Chama	Appro	1,232	19,712	2.6 GHz Intel Sandy Bridge:2S:8C	64	4	392	Sep-12
Uno	Dell	251	3,344	2.7 GHz Intel Sandy Bridge:2S:8C/4S:8C	64/128	4/8	71	Aug-14
Serrano	Penguin	1,152	41,472	2.2 GHz Intel Broadwell:2S:36C	128	3.5	1,382	Summer-16
Solo	Penguin	384	13,824	2.2 GHz Intel Broadwell:2S:36C	128	3.5	460	Summer-16
<b>Black Total:</b>		<b>7,690</b>	<b>130,504</b>				<b>3,157</b>	
Red Sky	Sun	519	4,152	2.93 GHz Intel Nehalem:2S:4C	24	3	48	Nov-10
Jemez	HP	288	4,608	2.6 GHz Intel Sandy Bridge:2S:8C	32	2	95	Jun-14
Pecos	Appro	1,232	19,712	2.6 GHz Intel Sandy Bridge:2S:8C	64	4	392	Sep-12
Cayenne	Penguin	1,152	41,472	2.6 GHz Intel Broadwell:2S:36C	128	3.5	1,382	Summer-16
<b>Red Total:</b>		<b>3,191</b>	<b>69,944</b>				<b>1,917</b>	
Red Mesa	Sun	1,920	15,360	2.93 GHz Intel Nehalem:2S:4C	12	1.5	180	May-10
<b>Green Total:</b>		<b>1,920</b>	<b>15,360</b>				<b>180</b>	
NSCC2	Appro	1,848	29,568	2.6 GHz Intel Sandy Bridge:2S:8C	64	4	588	Nov-12
NSCC3	Appro	924	14,784	2.6 GHz Intel Sandy Bridge:2S:8C	64	4	294	Dec-13
<b>Orange Total:</b>		<b>2,820</b>	<b>45,504</b>				<b>893</b>	
<b>TOTAL:</b>		<b>15,621</b>	<b>261,312</b>				<b>6,147</b>	

Black Multi-Cluster file systems:  
 6.2 PB Lustre (/gscratch+/fscratch)  
 1.3 PB GPFS (/gpfs1)  
 90 TB NAS (home and projects)

Red Multi-Cluster file systems:  
 6.9 PB Lustre (/hscratch+/iscratch)  
 2.9 PB GPFS (/gpfs1)  
 90 TB NAS (home and projects)

# Services Provided

- One queue per funding source
- Each queue has a fairshare target, users' usage history is secondary factor in scheduling
- Priority is granted on as-needed basis, users must request, provide justification and be approved by management council
- Don't use QOS at this time
- Legacy on older cluster of dedicated nodes that users can buy access to
  - Source of user complaints
  - Time sink for staff

# Team Resources

- 25-30 staff. Subdivided into teams for:
  - Help desk
  - Operations
  - Application Support
  - File Systems
  - Scheduling
  - Monitoring
  - System Architects
- Scheduling is part-time job with support from other staff helping part-time

# Reporting and Accountability

- Model is investment, not charging
- Monthly usage reports sent to funding source management
- Allocations revisited annually, or when new clusters are deployed
- Efforts to get funding groups to manage user priorities within their allocation have been unsuccessful
- Funding groups who are consistently overusing their allocation are asked to increase their investment
- Rarely engage with individual users

# User Support

- Separate user support team fields simple problem tickets and questions
- Issues that can't be handled by user support are escalated to the HPC team
- Web site with documentation and examples
- Sporadic user meetings - not well attended

# Future

- Moving towards shorter job wall clock limits in more open environments
- This year will decommission 3 old systems and deploy 3 new systems
- Prototype SLURM scheduling running in a test environment