




Lessons Learned When Building a Greenfield HPC Ecosystem

Andrew Keen


Michigan State University





Terminology


- High Performance Computing
- Greenfield- In contrast to brownfield






About iCER / HPCC

- ‘Cyber-Enabled’
- 300+ nodes (500+ soon)
- > 1 PB storage
- High speed networks, GPUs, Xeon Phi accelerators, Large memory
- Software!
- People!









There's more than FLOPS?

It's an ecosystem:

- Users!
- Compute
- Storage
- Physical Infrastructure
- Management Tools
- Policies
- Education
- Community








Startup

- Big SMP system from a Famous Name
 - RFP benchmarks looked great!
- Actual workloads...
- Didn't have adequate I/O bandwidth






There's more than FLOPS?

It's an ecosystem:


- Users!
- Compute
- Storage
- Physical Infrastructure
- Management Tools
- Policies
- Education
- Community





Storage


- Fast Storage
 - Lustre; 350 TB, 9 GB/s
- Safe Storage
 - ZFS
 - Data Integrity, Snapshots, Replication
 - Fast-ish
 - First TB free
 - \$175/TB/year- competitive with offline cloud storage!





Configuration Management


- Never manage your systems by hand
- Can manage appliances/devices
 - NetApps
 - Junipers
 - VMWare
- Puppet – GIT environments are cool





Availability


- HPC is not HA*
 - but that doesn't mean you can't avoid disruptive pain points
 - Outages can be disruptive
 - Build redundancy based on budget and tolerance for disruption





Cluster Management


- Hardware Management
 - IPMI
 - Firmware Updates and configuration!
 - Lock-in





Security


- Reflect resource's goals
- Environment
- Trusts matter





Use Outside Resources


- Organization and Community, resources
- Don't reinvent the wheel
- Gitlab!
- <http://gitlab.msu.edu>





Communication


- No one reads bulk email
- Few people read personal email
- Social Media?
- Ticketing / Issue Tracking is *critical*





Physical Concerns


- Lots of power, small space
- Whole Room vs. Spot
- Containment
 - Easy to prototype!
- Long term?





CyberInfrastructure days


- October 24-25
- Open to the MSU community to learn and collaborate about MSU, national CI resources
- <http://tech.msu.edu/CI-Days>



New Compute Cluster

- 2x Intel Xeon Ivy Bridge E5-2670v2 (2.5 GHz, 20 cores total)
- 500 GB HDD
- FDR (56 gigabit) network

Type	Mem (GB)	Accelerators	Total Performance (GigaFLOPS)	Cost
Base	64		400	\$3,805
Big Mem	256		400	\$5,339
Bigger Mem	512		400	~\$12,000
GPU	128	2x NVIDIA K20	2400	\$7,900
Phi	128	2x Phi 5110p	2400	\$9,043





Conclusion

Questions?

keenandr@msu.edu

<http://contact.icer.msu.edu>

