



# Current and In-Development HPCC Computational Resources

Andrew Keen

Institute for Cyber-Enabled Research

Michigan State University

# Bio

- Started operations with HPCC in 2005
- Snout-to-tail / Farm-to-table
- Current: Senior HPC Architect

# Formality: Low

- This is a very casual presentation

# An hour on this?

Year	Name	Description	ppn	Memory	Nodes	Total Cores	
	Cluster	Nodes	Cores (per node)	Memory (per node)	Interconnect	TFLOPS	1008
	intel07	128	8	8 GB	SDR Infiniband	9.4	256
	gfx10*	32	8	18 GB	Gigabit Ethernet	7.4	1528
	intel10	192	8	24 GB	QDR Infiniband	14.7	64
	intel11	2	64	2 TB	QDR Infiniband	2.1	32
		1	32	1 TB	QDR Infiniband		128
		2	32	512 GB	QDR Infiniband		2560
	intel14	128	20	64GB	FDR Infiniband	51	480
		24	20	256GB	FDR Infiniband	9.6	800
	intel14-k20	40	20 (cpu)	128GB (host)	FDR Infinigand	108	560
			4992 stream processors	10GB (gpu)			240
	intel14-phi	28	20 (cpu)	128GB (host)	FDR Infiniband	67.2	96
			120 (phi)	16GB (phi)			
	<b>Total</b>					<b>580</b>	<b>7752</b>



# What I'm Talking About

- Compute
- Storage
- Network
- Services\*

# What I'm Talking About

- Past
  - **“The past is never dead. It's not even past.”**
- Present

# First Systems (2005)

- Green
  - 256-576-512 GB, 64-128-112 cores.
  - Itanium2

m  
ems



sgi

sgi

sgi allix 3000



# AMD05

- Shaw (AMD05)
  - 128 nodes
  - 8 GB, 4 cores, SDR Infiniband
  - 4 DP OPs/clock
  - 2.4 TF

# Lustre06

- 6 TB, 1 GB/s, 10K FC back-end

# Intel07

- SGI
- 128 nodes, 8 cores, 8 GB, SDR  
Infiniband

# Intel08

- 16 core, 64 GB x2

# Lustre09

- 100 TB
- 2 GB/s
- SATA
- Sun Thumper

# AMD09

- 32 cores, 256 GB x5

# Legacy

# Intel10

- 192 nodes
- 24 GB
- 8 cores
- QDR Infiniband
- 4 DP OPS/clock
- 15 TF



# GFX10

- 32 nodes, 2x GPUs (80 GF each)

# Current

# intel11

- 5 nodes, 512 GB to 2 TB, 32-64 cores, QDR IB.

# intel14

- 128 20 core, 64 (or more)
- 24 20 core, 256 GB

# intel14-GPU

- 40x 2x K20

# intel14-Phi

- 28x 2x Phi 5110

# intel14-XL

- 6 nodes, 1-6 TB, 48-96 3GHz IVB cores

# 'eval' stack

- One-offs, misfits, temporary machines
- Previously: GPUs, microservers, Cell
- Currently: eval-k10, eval-k40



# Storage

- Home
- Scratch
- OSIRIS

# Home Directories

- About 1 PB usable, SSD-accelerated
- Replicated
- Snapshots

# Home Directories

- Users can access snapshots:
- `$USER/.zfs`
- Research spaces are available for collaborative work

# Lustre '12

- 350 TB, 10 GB/s

# Lustre '15

- 1.9 PB
- 35 GB/s
- FDR IB (56 gigabit) fabric
- POSIX, ACLs, ZFS backend
- Faster metadata, small file performance

# OSIRIS

- MI-OSIRIS, NSF-funded distributed software-defined storage platform

# OSIRIS

- Multidisciplinary, multi-institution, multiprotocol

# OSIRIS

- Storage Building Block



# OSIRIS

- Network Infrastructure

# OSIRIS

- Application, Metadata

# Networking

- Core network at 56 gigabit
- Low-latency
- Fabric unification planned

# External Networking

- Upgrades possible
- MSU campus networking review underway

# Compute Services

# Traditional HPC

- Moab for Workload Management
- TORQUE for Job Management

# Software Environment

- Maintained research application platforms

# System queueing model

- Condoized Library
- Fair-share
- Single\* Queue



# Gateways

- RDP Gateway
- Application Gateways
- VCL

# Questions?