

Hyak Governance Board Future Overview

November 13, 2015


Agenda

- Hyak Classic
 - Timeline
 - Budget
- NextGen Hyak
 - Design
 - Timeline

Hyak Classic Retirement Timeline

- Committed to run through FY20
- More or less steady state through FY18
- A&S, CoE, HHMI, and UW-IT = large commitments

Hyak Classic Node Retirement Schedule -- LOW ESTIMATE



Hyak Classic Node Retirement Schedule -- HIGH ESTIMATE









Chart - Node Timeline by Sponsor - HIGH



Hyak Classic Budget Estimates

\$0 \$50,000 \$100,000 \$150,000 \$200,000 \$250,000 \$300,000 \$350,000 \$400,000 \$450,000



■ High Budget ■ Low Budget

NextGen Hyak Design

- Interconnect
- CPU Support
- Accelerator Support
- Scratch Storage
- Scheduler
- Timeline
- Design Patterns
- Budget

NextGen Hyak Interconnect

- Ethernet
 - Not ready for prime time
 - Too expensive
- Infiniband
 - Industry standard today
 - Relatively expensive
- **Intel Omnipath**
 - **New, but HUGE baking**
 - **Superfast (100Gbs, < 150ns)**
 - **Good value**

NextGen Hyak CPU

- IBM Power
 - New
 - Fast
 - Semi-open
 - Used on some new DoE systems
 - Too expensive
- **Intel Xeon**
 - **Industry standard**
 - **Fast**
 - **Most open**
 - **Most widely used**
 - **Best value**

NextGen Hyak Accelerators

- NVIDIA
 - Industry standard
 - Slow adoption
 - Few real world successes
 - Used on most large DoE & NSF systems
 - Too expensive
- Intel Xeon Phi (Knight's Corner/Landing)
 - New
 - Anticipate faster adoption
 - Too new to know how things will unfold
 - Used on an increasing number of NSF & DoE systems
 - Value unknown

NextGen Hyak Scratch

- > 100GB?
- > 500GB?
- > 1,000GB?

NextGen Hyak Scheduler

- Hyak Classic uses Moab/Torque (PBS)
 - Was Industry standard in 2010
 - Expensive
 - Complex
 - Poor support
- NextGen Hyak will use SLURM
 - Emerging as industry standard
 - Open source Good design
 - Inexpensive support

NextGen Hyak Timeline

- Technology Timeline
 - 4/1/2016 - Lenovo BCH Slowdown
 - 7/1/2016 - Lenovo BCH EOL
 - Q1/Q2 2016 - Intel Omnipath GA
 - Q1/Q2 2016 - Intel Knight's Landing GA

NextGen Hyak Timeline

- Managing the Gaps
 - Minimize the Classic -> NextGen Gap
 - Match our timeline to vendor's (OP & KL)
 - Match our timeline to grants requiring match (NSF MRI)

NextGen Hyak Timeline

- Design Patterns
 - Deploy network backbone and anchor tenants
7/1/2014 (~100 nodes)
 - Scale out one rack at a time (~72 nodes)
 - Max out at ~1,000 nodes
 - Node lifetime four years
 - Three year build + four year run, decommission —
Seven years total

NextGen Hyak Timeline

- **SPONSORS PLEASE SURVEY YOUR USERS!!!**