

THE TEXAS ADVANCED COMPUTING CENTER TACC IN 2017

Dan Stanzione Executive Director Austin, TX dan@tacc.utexas.edu



6/20/17

STAMPEDE 2 AT TACC

The Xeon Phi (Knights Landing) component of Stampede-2 has just begun production (all users added one week ago).

- We began experimenting with a smaller number of KNLs this time last year as part of an upgrade to Stampede-1.
- Stampede-1 included KNC cards when it launched in early 2013, so we are approaching five years with Xeon Phi.



STAMPEDE 2 SPECS

▶ Phase 1: 4,200 Intel Xeon Phi 7250 Nodes produced by Dell.

- ► ~30PB Seagate scratch filesystem
- ▶ (Now in service!)

Phase 2: 1,736 Intel Xeon Sky Lake nodes from Dell – later this summer – shares one Omnipath Fabric with Phase 1.

Phase 3: A small, experimental, 3D Xpoint (DIMM) upgrade in 2018.



EARLY RESULTS

- With the first review coming up, most will be announced in the next few weeks, but a few things I can say now:
 - ▶ 100% of our acceptance tests passed
 - We have had some *great* results with highly tuned Seismic codes that take full advantage of the HBM.
 - Some other "newer" codes have also done great several early user papers just accepted for SC17, so check back in the fall.
 - Most of our "major" applications are at 2-3x Stampede 1 per node performance
 - Recently ported in the Gordon Bell winning code from 2015 and a finalist from 2013, with good results.



EARLY RESULTS - 2

▶ If I had to generalize:

Everything runs, but...

- Carefully tuned codes are doing pretty well, but with work.
- "Traditional" MPI codes, especially with OpenMP in it do relatively well, but not great.
- Some codes, particularly, not very parallel ones, are pretty slow, and probably best off on Xeon.

Mileage varies widely – but raw performance isn't the whole story.



MY PERSONAL OPINION ON XEON PHI

Xeon Phi looks to be the most cost and power efficient way to deliver performance to highly parallel codes.

- In many cases, it will not be the fastest. For things that only scale to a few threads, it is *definitely* not the fastest.
- But what is under-discussed:
 - As far as I can tell, a Dual-socket Xeon node costs 1.6x what a KNL node costs, even after discounts.
 - ► A dual-socket, dual GPU nodes is probably >3x a Xeon Phi node.
 - A KNL node uses about 100 less watts per node than a dual-socket Xeon node.
- Is "fastest" our only metric? We seem to be heading into a very workload-dependent age when discussing best value.



THANKS!

DAN STANZIONE DAN@TACC.UTEXAS.EDU

