




OMNI-PATH PANEL

SC17

Dan Stanzione, Executive Director

Austin, TX

dan@tacc.utexas.edu



Follow the
STAMPEDE

Powering Discoveries That Change The World

STAMPEDE 2

- ▶ Funded by NSF as a renewal of the original Stampede project.
- ▶ The largest XSEDE resource (and largest university-based system).
- ▶ Follow the legacy of success of the first machine as a supercomputer for a *broad* range of workloads, large and small.
- ▶ Install without ever having a break in service – in the same footprint.



STAMPEDE 2 -- COMPONENTS

▶ Phase 1

- ▶ 4,204 Intel Xeon Phi "Knights Landing" (KNL) Processors (Intel and Dell)
- ▶ ~20PB (usable) Lustre Filesystem (Seagate), 310GB/s to scratch.
- ▶ Intel Omni-Path Architecture fabric – Fat Tree.
- ▶ Ethernet fabric and (some) management infrastructure.

▶ Phase 2

- ▶ 1,736 Intel Xeon (Sky Lake) Processors
- ▶ (Associated networking, but core in phase 1).
- ▶ Balance of management hardware.

▶ Phase 3

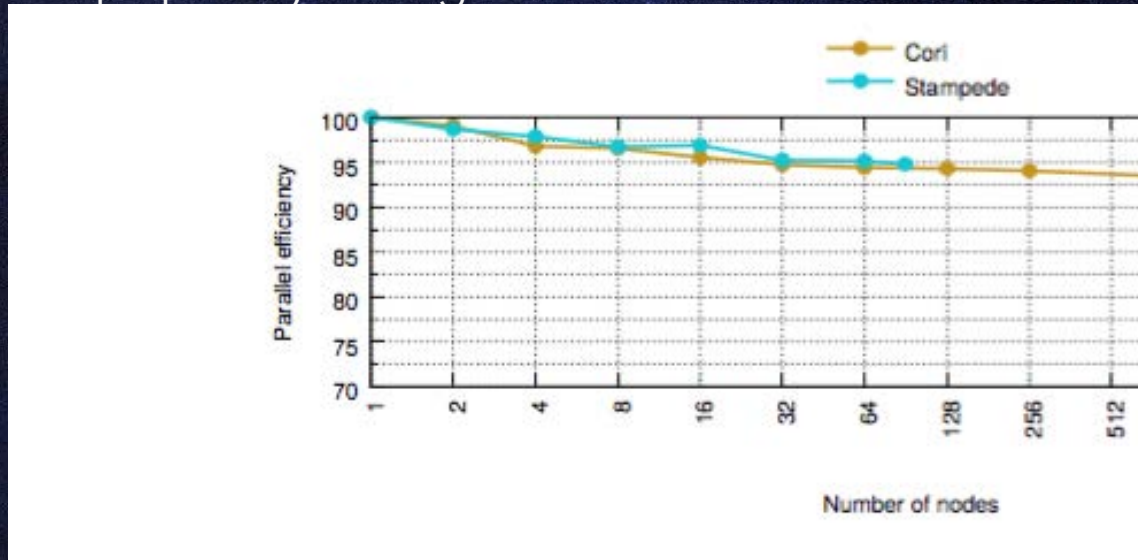
- ▶ Intel 3D XPoint DIMMS as an experimental component in a small subset of the system.

OPA FABRIC TOPOLOGY

- ▶ Fat-tree topology design with 7:5 oversubscription
- ▶ Each top of rack switch connects to twenty different line cards to flatten topology
 - ▶ Up to sixteen switches per line card
 - ▶ Switches in the same rack always connect to same line card
- ▶ Adjacent racks connected to same line card as much as possible
 - ▶ E.g. first 8 racks connect to same line card
- ▶ Intel tested system topology with their simulator
- ▶ I/O switches spread across line cards to avoid I/O bottlenecks
- ▶ Custom cable management panels to allow for easy cabling of core switches

OMNIPATH SCALING

- ▶ Seems just fine
 - ▶ comparison to KNL-Cray at NERSC, just to 128 nodes (on a seismic code).
 - ▶ From an ISC paper by Yifeng Cui



THANKS!

DAN STANZIONE
DAN@TACC.UTEXAS.EDU