



# STAMPEDE 2 UPDATE

John Cazes, Texas Advanced Computing Center



TEXAS



*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 in the US)
- Follows the legacy of success of Stampede 1 as a supercomputer for a \*broad\* range of workloads, large and small
- Installed without ever having a break in service – in the same footprint
- Composed of Intel “Knights Landing” Xeon Phi (Phase 1) and Intel “Sky Lake” Xeon (Phase 2) processors; 18PF total



# STAMPEDE 2 – 3 PHASES

- Phase 1 – June 2017
  - 4,204 Intel Xeon Phi 7250 "Knights Landing" (KNL) nodes
  - ~20PB (usable) Lustre Filesystem (Cray), 310GB/s to /scratch.
  - Intel OmniPath Architecture (OPA) Fabric – Fat Tree topology
  - Ethernet fabric and (some) management infrastructure.
- Phase 2 – December 2017
  - 1,736 Intel Xeon Platinum 8160 "Skylake" two-socket nodes
  - (Associated rack level networking, but core in phase 1).
  - Balance of management hardware, new Skylake servers
- Phase 3 – 2<sup>nd</sup> half 2018
  - 3D Xpoint NVDIMMS as an experimental component in a small subset of the system.

# TOP 500 – NOVEMBER 2017

## Stampede 2 comes in at #12

Rank	System	Cores	Rmax (TFlop/s)	Rpeak (TFlop/s)	Power (kW)
12	Stampede2 - PowerEdge C6320P/C6420, Intel Xeon Phi 7250 68C 1.4GHz/Platinum 8160, Intel Omni-Path , Dell EMC Texas Advanced Computing Center/Univ. of Texas United States	368,928	8,317.7	18,215.8	

Worked on improving our percentage of peak

- Binned nodes in terms of peak HPL performance
- Changed the amount of work/node to match performance bin
- Enabled HUGE pages

# TOP 500 – JUNE 2018

Performance improves by over 2 petaflops!

- Changing the way we run HPL paid off
- Credit goes to Tommy Minyard & John McCalpin

Rank	System	Cores	Rmax (TFlop/s)	Rpeak (TFlop/s)	Power (kW)
15	Stampede2 - PowerEdge C6320P/C6420, Intel Xeon Phi 7250 68C 1.4GHz/Platinum 8160, Intel Omni-Path , Dell EMC Texas Advanced Computing Center/Univ. of Texas United States	367,024	10,680.7	18,309.2	

# TOP 500 – JUNE 2018

Performance improves by over 2 petaflops!

- Changing the way we run HPL paid off
- Credit goes to Tommy Minyard & John McCalpin

Rank	System	Cores	Rmax (TFlop/s)	Rpeak (TFlop/s)	Power (kW)
15	Stampede2 - PowerEdge C6320P/C6420, Intel Xeon Phi 7250 68C 1.4GHz/Platinum 8160, Intel Omni-Path , Dell EMC Texas Advanced Computing Center/Univ. of Texas United States	367,024	10,680.7	18,309.2	


Although 2PF faster, we slide from 12 to 15

# SCIENCE HIGHLIGHTS

# REAPING POWER FROM WIND FARMS

## Multi-Scale Model of Wind Turbines

- Optimized control algorithm improves design choices
- New high-res models add nacelle and tower effects
- Blind comparisons to wind tunnel data demonstrate dramatic improvements in accuracy
- Potential to increase power by 6-7% (\$600m/yr nationwide)



Graphic from Wind Energy, 2017.

*"TACC...give[s] us a competitive advantage..."*

Christian Santoni, Kenneth Carrasquillo,  
Isnardo Arenas-Navarro, and Stefano Leonardi

UT Dallas, US/European collaboration (UTRC, NSF-PIRE 1243482)

For more details




*"Using commodity HPC servers...the time to data-driven discovery is reduced and overall efficiency can be significantly increased." (Niall Gaffney, TACC)*

## RECORD ACHIEVED ON AI BENCHMARK

TACC, Berkeley, Cal Davis collaborate on large-scale AI runs

- Research demonstrating the potential of commodity hardware for AI
- Skylake ImageNet benchmark: (100 epochs, 11 min, 1024 nodes) -- fastest result at time of publication
- Knights Landing ImageNet benchmark (90 epochs, 20 min, 2048 nodes) – 3x faster than Facebook, with higher large-batch accuracy

Graphic credit Andrej Karpathy




Yang You, Zhao Zhang, Cho-Jui Hsieh, James Demmel, Kurt Keutzer

## USING KNL TO PROBE SPACE ODDITIES

*Ongoing XSEDE collaboration focusing on KNL performance for new, high-resolution version of COSMOS MHD code*




- Vectorization and other serial optimizations improved KNL performance by 50%
- COSMOS currently running 60% faster on KNL than Stampede1
- Work on OpenMP-MPI hybrid optimizations now underway
- Impact of performance improvements amounts to millions of core-hours saved



*"The science that I do wouldn't be possible without resources like [Stampede2]...resources that certainly a small institution like mine could never support. The fact that we have these national-level resources enables a huge amount of science that just wouldn't get done otherwise." (Chris Fragile)*

XSEDE ECSS: Collaboration between PI Chris Fragile (College of Charleston) and Damon McDougall (TACC)

# RESPONDING TO HURRICANE HARVEY



- Next Generation Storm Forecasting (with Penn State)
- Storm Surge Modeling (with Clint Dawson UT Austin)
- Preliminary river flooding and inundation maps (David Maidment UT Austin)
- Remote Image Integration and Assimilation (Center for Space Research, UT Austin)

**THANKS!**

John Cazes  
cazes@tacc.utexas.edu