

# Argonne Training Program on Extreme-Scale Computing (ATPESC)

## Introduction to ATPESC



**Michael E. Papka**  
Argonne Leadership Computing Facility  
Argonne National Laboratory




Q Center, St. Charles, IL (USA)  
July 29, 2019



# Filling in for Pete Beckman



# Watch Pete's Videos

|  |  |
|--|--|
|   | <p><b>An Introduction to Parallel Supercomputing   Pete Beckman, Argonne National Laboratory</b></p> <p>ANL Training • 359 views • 8 months ago</p> <p>Presented at the Argonne Training Program on Extreme-Scale Computing 2018. Slides for this presentation are available here: ...</p> |
|   | <p><b>An Introduction to Parallel Supercomputing   Pete Beckman, Argonne National Laboratory</b></p> <p>ANL Training • 1.8K views • 2 years ago</p> <p>Presented at the Argonne Training Program on Extreme-Scale Computing, Summer 2016. Slides for this presentation are ...</p>         |
|  | <p><b>An Introduction to Parallel Supercomputing   Pete Beckman, Argonne National Laboratory</b></p> <p>ANL Training • 448 views • 1 year ago</p> <p>Presented at the Argonne Training Program on Extreme-Scale Computing 2017. Slides for this presentation are available here: ...</p>   |

# Pete's Advice

- Put phones away
- Ignore email for a few weeks
- Dive into the material presented / experiment
- Have fun ...

# Agenda

# Argonne First Digital Computer




AVIDAC, Argonne's first digital computer, began operation in January 1953. It was built by the Physics Division for \$250,000.

# Established Advanced Research Computing Facility (ARCF)



# Continued Deployment of Computing Advances





# Fast Forward to Today




**Mira IBM BG/Q**  
 49,152 nodes  
 786,432 cores  
 768 TiB RAM  
 Peak flop rate: 10 PF

**Cetus IBM BG/Q**  
 4,096 nodes  
 65,536 cores  
 64 TiB RAM  
 Peak flop rate: 836 TF

## Tape

- The ALCF has three 10,000-slot libraries using LTO 6 tape technology.
- Hardware compression for an effective capacity of 36-60 PB.



**Theta Cray XC40**  
 4,392 nodes  
 281,088 cores  
 892 TiB RAM  
 Peak flop rate: 11.69 PF

**Iota Intel/Cray XC40**  
 44 nodes  
 2,816 cores  
 8.9 TiB RAM  
 Peak flop rate: 117 TF

## Storage Capability

### Disk

- Mira: ~27 PB of GPFS file system capacity with performance of 240 GB/s on the largest file system (19PB).
- Theta: ~18 PB of GPFS/Lustre file system capacity; 9PB is GPFS and 9.2PB is Lustre.

**Cooley Cray/NVIDIA**  
 126 nodes  
 1512 Intel Haswell CPU cores  
 126 NVIDIA Tesla K80 GPUs  
 48 TB RAM / 3 TB GPU

**Firestone IBM Power8**  
 2 nodes + K80 GPU  
 20 cores  
 128 GB RAM  
*Hybrid CPU/GPU*

# High Level View of Aurora

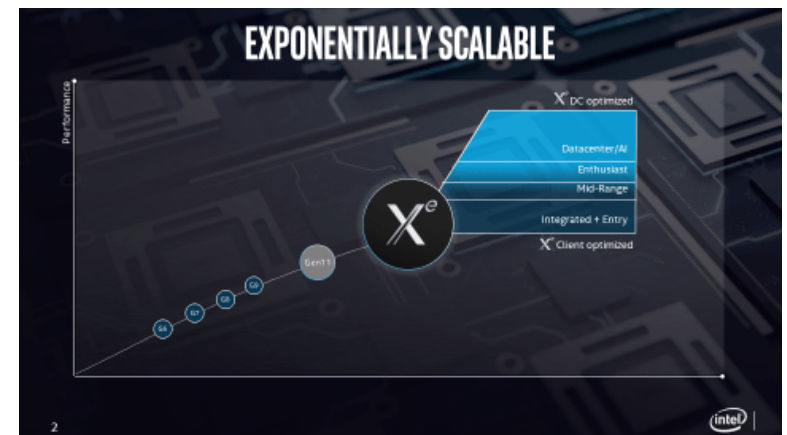
- Aurora is an Intel/Cray machine
- The system is to be delivered to Argonne in 2021
  - Limited availability to applications in 2021
  - Expect broader availability to applications in 2022
- Aurora will be a Exa-scale system
  - Will have a peak performance of over 1 ExaFlop/s in Double-Precision
  - Much higher than 1 ExaFlop/s in Half-Precision
  - Target for applications performance is an 50x speedup over Titan/Sequoia



# Aurora Hardware Overview

## Compute Node and Memory

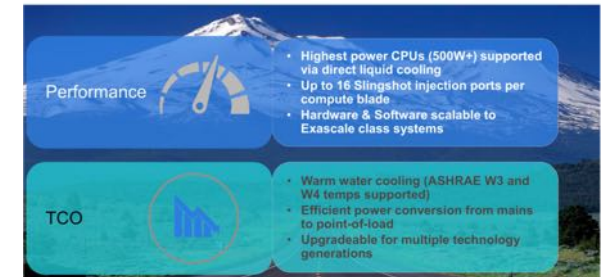
- Processor
  - Future Intel® Xeon® Scalable Processor
  
- Accelerator
  - New Intel® X<sup>e</sup> Compute Architecture



# Aurora Hardware Overview

## Platform, Fabric and I/O

- Compute Platform
  - Cray Shasta next generation supercomputing platform
- System Interconnect
  - Cray Slingshot interconnect
- I/O system
  - Intel's DAOS (Distributed Application Object Storage)
  - Traditional parallel filesystem augments DAOS (bulk storage, legacy support)



# Top500.org



# QUESTIONS?

Michael E. Papka  
papka@anl.gov