Next Supercomputer at CCS: Cygnus-BD – Big memory supercomputer

Our targets

- Accelerates large-scale data analysis and big data AI by utilizing persistent memory for large memory space and high performance storage
- Fosters new fields of large-scale data analysis, new applications of big data AI, and system software research

Background

- In eight years, **CPU performance increased 50 times**, while **memory performance increased** only **3.8 times**
- Memory size and storage performance matter

Persistent Memory

- One order better cost performance
- Minimum latency is ~60 ns (similar to DRAM)
- Half of bandwidth
- Memory mode
  - Larger memory space without much performance penalty
- Direct mode
  - Direct access to byte-addressable persistent memory and high-performance storage

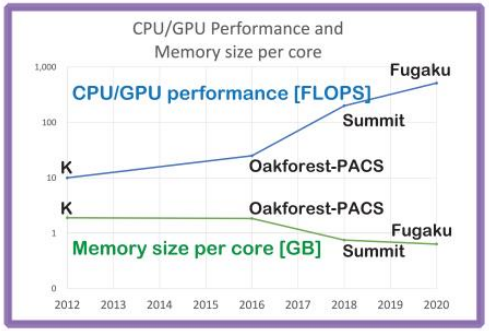

Cygnus-BD estimated performance

- Total Performance
  - 300 nodes, 7 PFlops, 340 TB
- Node performance
  - 3 TFlops (CPU), 20 TFlops (GPU)
  - 128 GB DRAM, 2 TB Pmem
  - 3 TB NVMe SSD
- Interconnection Network
  - 200 Gbps full bisection
- Parallel File System
  - 7 PByte, 50 GB/s

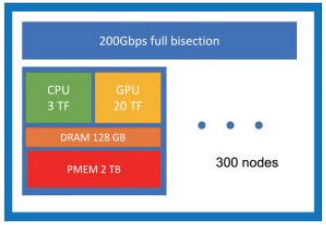

	Cygnus-EC (2019)	Cygnus-BD (2022)
PFLOPS (DP)	2.4	6.9
CPU	0.16	0.9
GPU	2.24	6.0
Memory (TiB)	15.2	37.5
Pmem (TiB)	0	600
Storage (PB)	2.4	7

Schedule

2021 Q1	Q2	Q3	Q4	2022 Q1	Q2	Q3	...
RFI/RFC		RFP		Installation		Operation	

CHFS ad hoc distributed file system

- Temporal distributed file system using node-local storage
  - Aka node-local burst buffer
- Fill the performance gap between CPU/GPU and storage

- We are developing CHFS ad hoc file system to utilize persistent memory
  - No metadata server, no sequential processing for performance and scalability

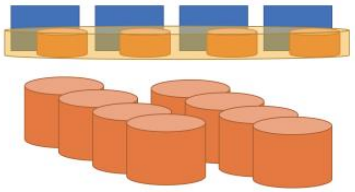

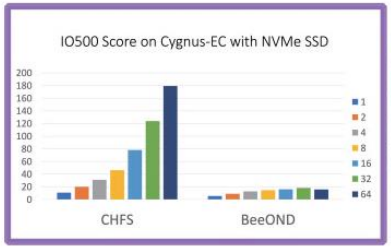

Cygnus-EC