



Blue Waters Supercomputer



- NSF-funded (\$200M) Supercomputer @ U. Illinois & NCSA
- Cray XE6/XK7 hybrid machine composed of
 - AMD 6276 "Interlagos" processors
 - NVIDIA GK110 "Kepler" accelerators (XK nodes)
 - Cray Gemini 3D Torus interconnect
- **Compute**
 - 23k Cray XE nodes
 - 362k Bulldozer cores; 1.4 PB memory
 - 4.2k Cray XK nodes
 - 34k Bulldozer cores + 4.2k Kepler accelerators
 - 135 TB CPU / 25 TB GPU memory
- **Storage, I/O**
 - Online: 26 PB, aggregate I/O > 1 TB/sec
 - Near-line: 380 PB, 1.2 GB cache front-end, 58 GB/s aggregate to tape



National Petascale Computing Facility





Blue Waters Supercomputer



XE Compute Node

AMD 6276 Interlagos Processors	2
Bulldozer Cores*	16
Integer Scheduling Units**	32
Memory / Bulldozer Core	4 GB
Total Node Memory	64 GB
Peak Performance	313.6 GF
Memory Bandwidth	102.4 GB/s

XK Compute Node

AMD 6276 Interlagos Processors	1
Bulldozer Cores*	8
Integer Scheduling Units**	16
Memory / Bulldozer Core	4 GB
Node System Memory	32 GB
GPU Memory	6 GB
Peak CPU Performance	156.8 GF
CPU Memory Bandwidth	51.2 GB/s
CUDA cores	2688
Peak GPU Performance (DP)	1.31 TF
GPU Memory Bandwidth (ECC off)***	250 GB/s

Interconnect

Architecture	3D Torus
Topology	24x24x24
Compute nodes per Gemini	2
Peak Node Injection Bandwidth	9.6 GB/s

Online Storage

Total Usable Storage	26.4 PB	
Aggregate I/O Bandwidth	> 1 TB/s	
File System	Size (PB)	# of OSTs
home	2.2	144
projects	2.2	144
scratch	22	1440

See: <https://bluewaters.ncsa.illinois.edu>



Blue Waters Supercomputer



- **Policy:** “2% of the available time will be allocated to university projects: (i) faculty whose research and/or education programs would be greatly enhanced by access to Blue Waters and (ii) research and/or education proposals where a commitment of Blue Waters resources will significantly increase the competitiveness of the proposals”
- **Three types of allocations: Exploratory, General, Education**
- **General proposal** (Intended for large-scale research projects)
 - **March 2016:** submitted a proposal (PI: Neubauer) for **1M** node-hours
 - **May 2016:** Awarded **50k** node-hours, **5TB/50/500TB** for home/project/scratch managed by Cray’s Sonexion Lustre system
 - ❖ “Blue Waters (BW) Project Office has requested an initial award to explore the technical feasibility of your proposal. Your full request is not rejected, but rather deferred for later review pending this initial technical feasibility...”
 - **August 1, 2016:** Start date for allocation. Met last week with BW technical team and developed plan for initial deployment of services



Technical Plans

- BW Scheduler Submission
- Pilot Submission
 - Internal APF/RCCF deployment on edge device
 - External APF/RCCF using gsiftp and OTP (weekly update)
- Software Deployment
 - Shifter container with complete Atlas installation
 - Minimal node with fuse and VM via CVMFS
 - Without fuse, Stratum-R availability of containers
- Deploy RSE
 - Initially use remote StageIn/Out of data via limited NAT
 - Pursue RSE edge device to cache local data
 - Use Globus FTP servers to transfer data in/out
- Configure `CONNECT_BLUEWATERS`
 - If using RSE, LSM to find data