


BLUE WATERS

SUSTAINED PETASCALE COMPUTING

Blue Waters System Overview


Greg Bauer




GREAT LAKES CONSORTIUM
FOR PETASCALE COMPUTATION

CRAY®

The Blue Waters EcoSystem



Blue Waters Computing System



Aggregate Memory – 1.5 PB

Scuba Subsystem -
Storage Configuration
for User Best Access

120+ Gb/sec

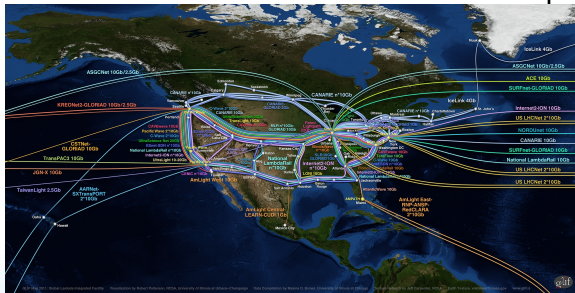
10/40/100 Gb
Ethernet Switch

External Servers

IB Switch

>1 TB/sec

100 GB/sec



100-300 Gbps WAN



Spectra Logic: 300 usable PB



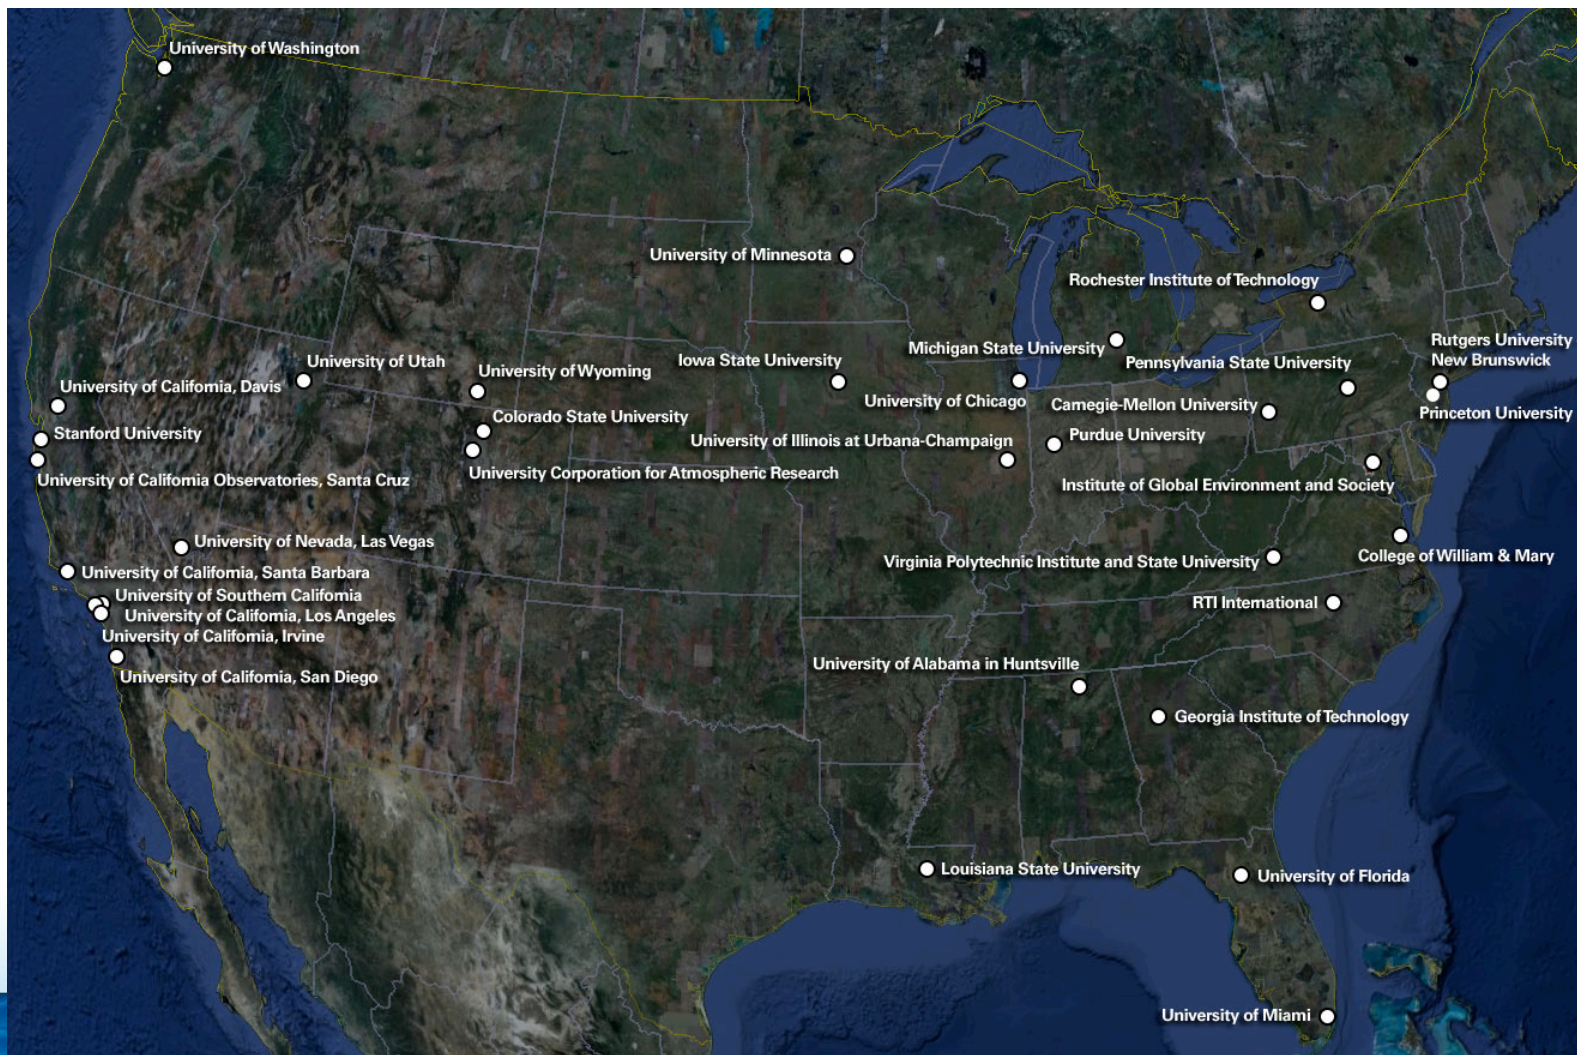
Sonexion: 26 usable PB

National Petascale Computing Facility

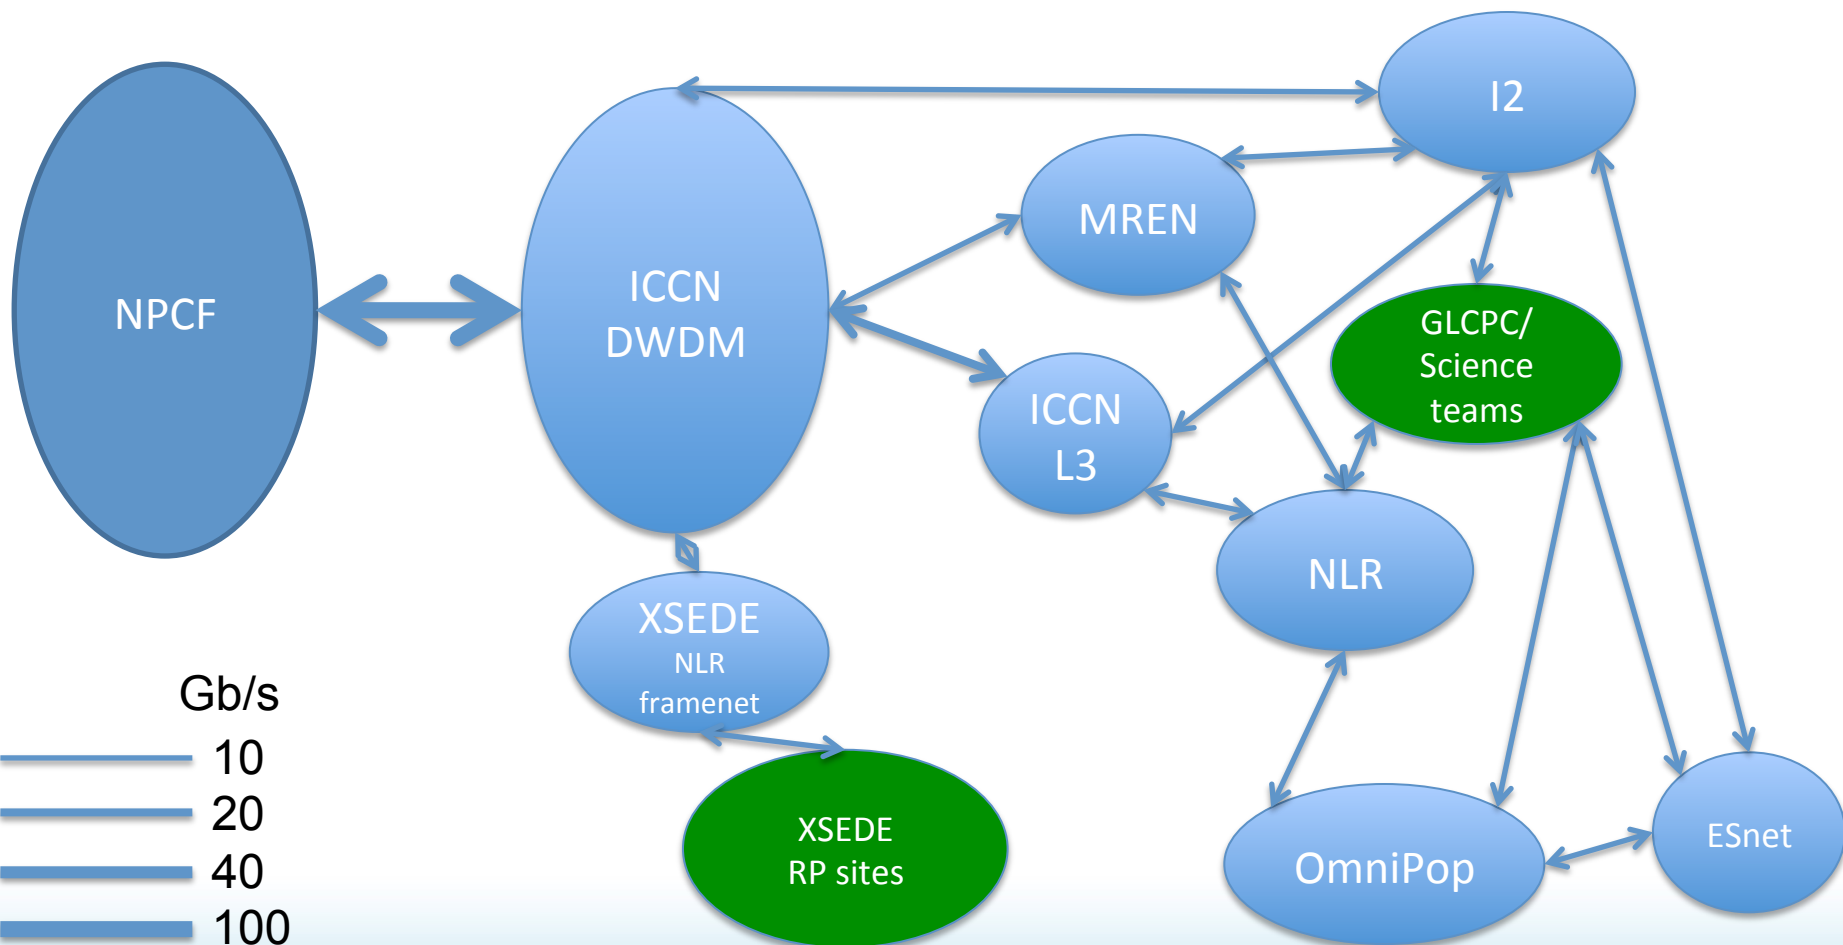


- Modern Data Center
 - 90,000+ ft² total
 - 30,000 ft² 6 foot raised floor
 - 20,000 ft² machine room gallery with no obstructions or structural support elements
- Energy Efficiency
 - LEED certified Gold
 - Power Utilization Efficiency, PUE = 1.1–1.2
 - 24 MW current capacity – expandable
 - Highly instrumented

PRAC PI Institutions – June 2012

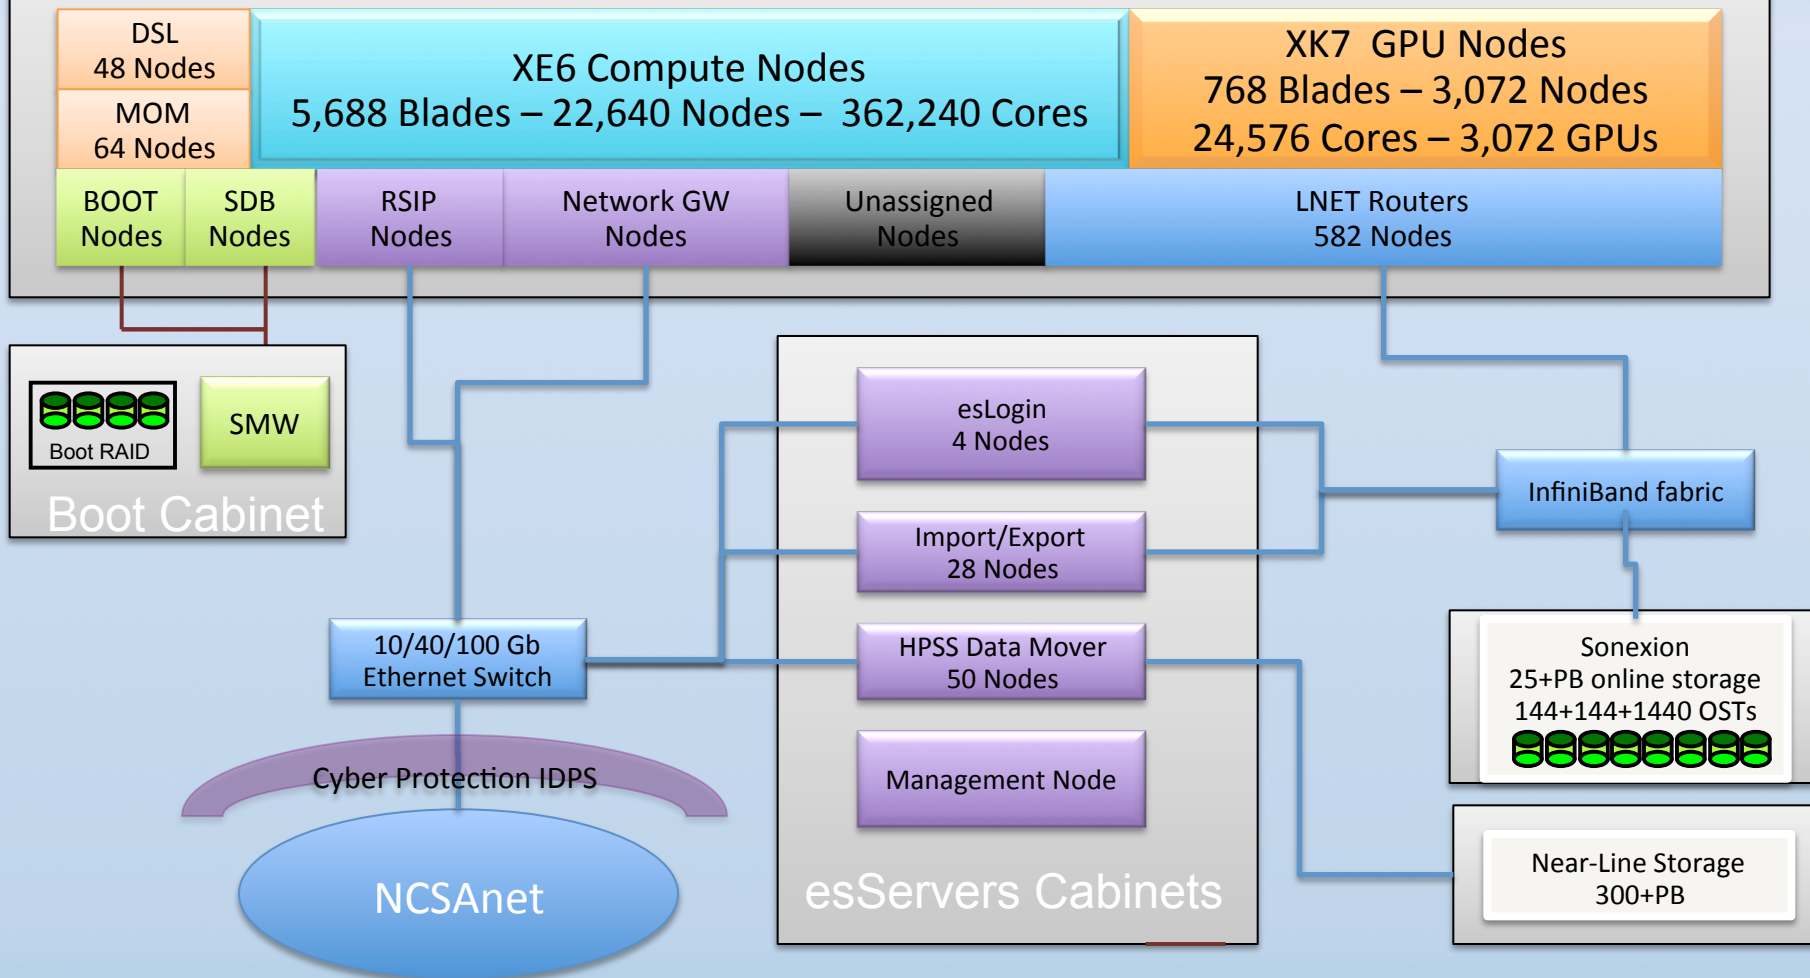


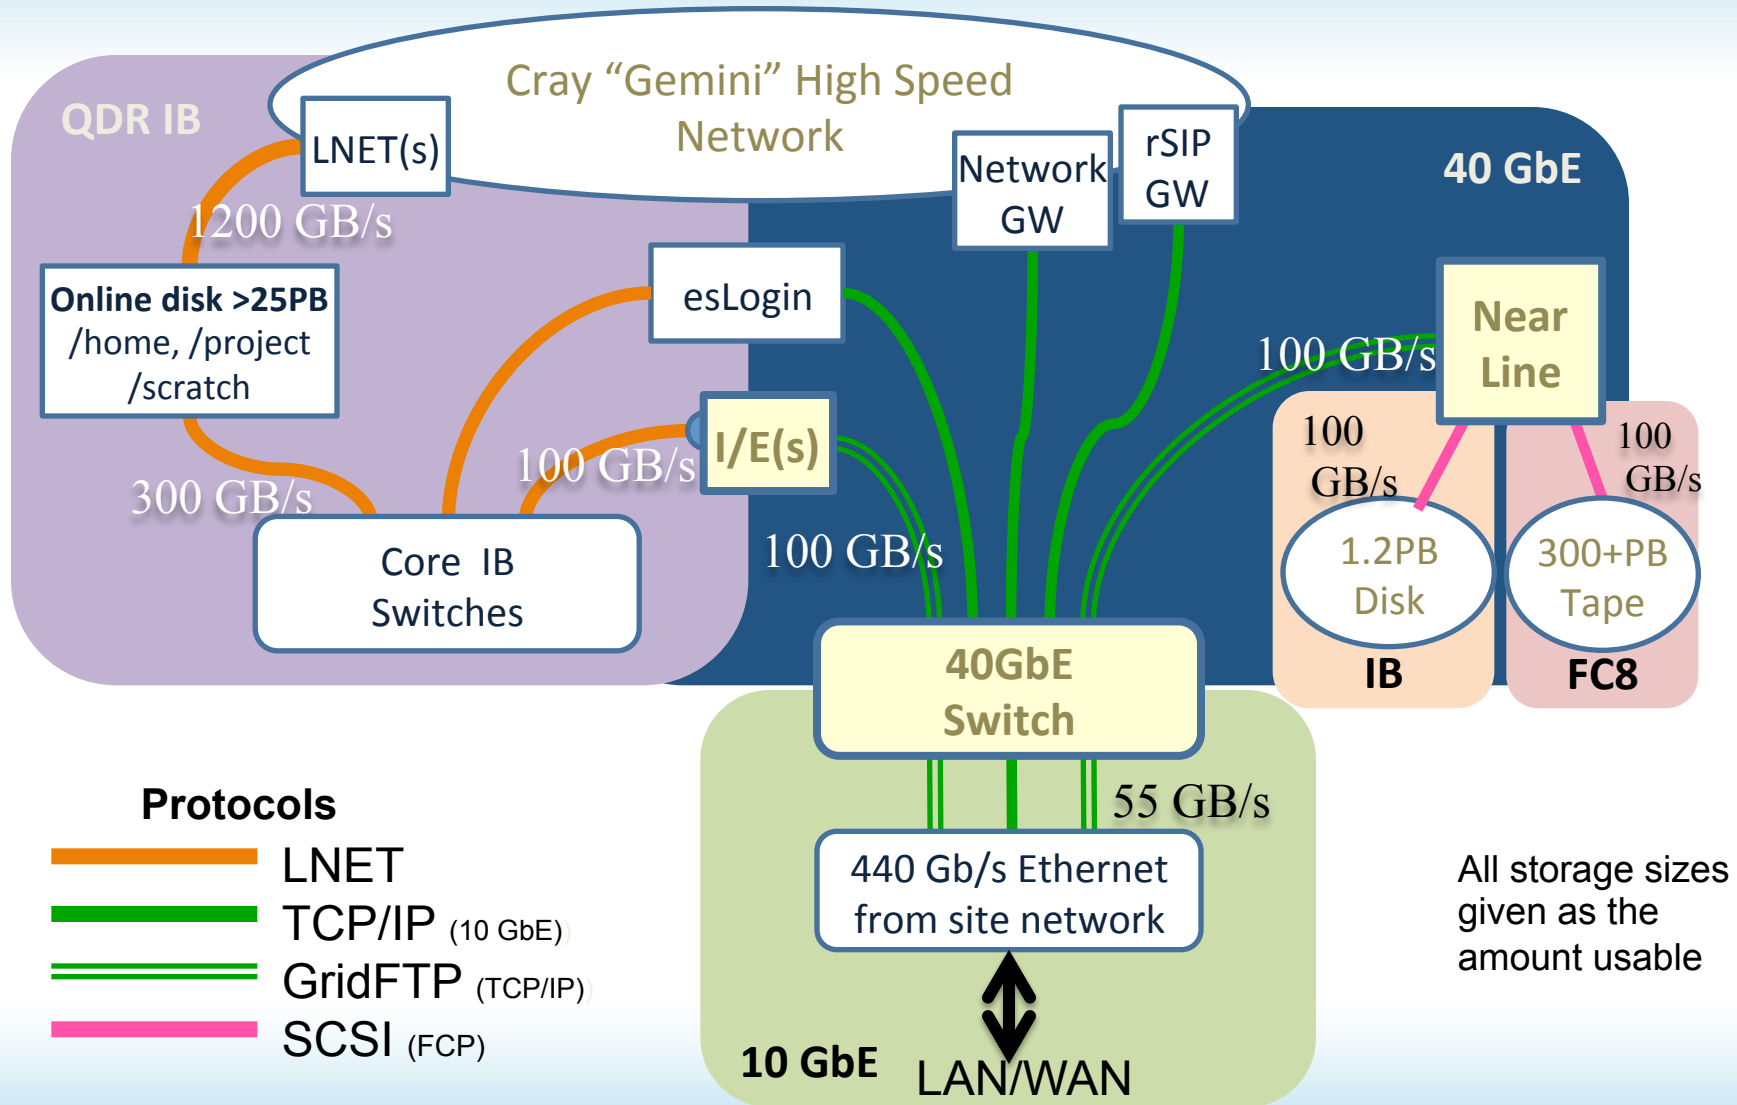
The Movement of Data



Gemini Fabric (HSN)

Cray XE6/XK7 - 276 Cabinets



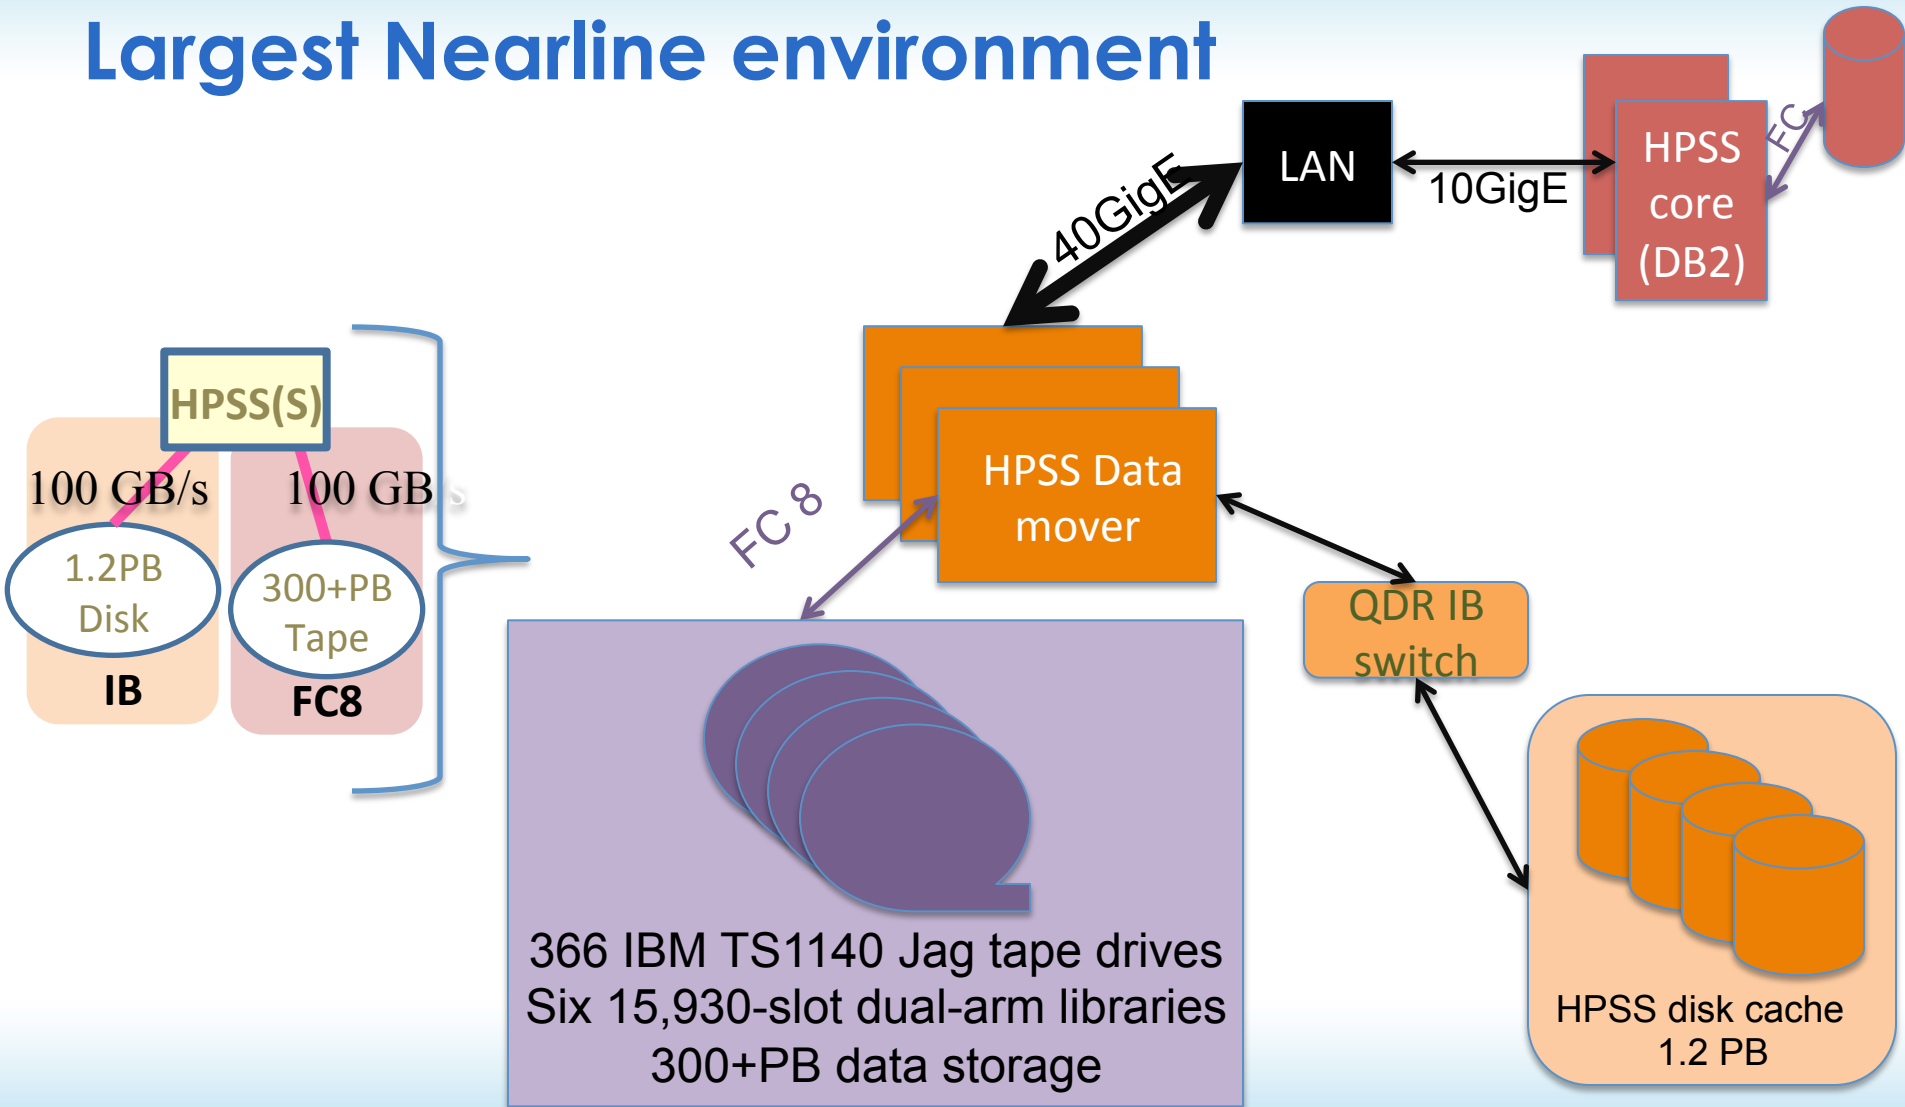


Blue Waters Nearline/Archive System

- Spectra Logic T-Finity
 - Dual-arm robotic tape libraries
 - High availability and reliability, with built-in redundancy
- Blue Waters Archive
 - Capacity: 380 PBs (*raw*), 300 PBs (*usable*)
 - Bandwidth: 100 GB/sec (*sustained*)
 - Redundant arrays of independent tapes RAIT for increased reliability



Largest Nearline environment



Online Storage



home : 144 OSTs : 2.2 PB useable : 1 TB quota



projects: 144 OSTs : 2.2 PB useable : 3 TB group quota



scratch: 1440 OSTs : 22 PB useable : 500 TB group quota

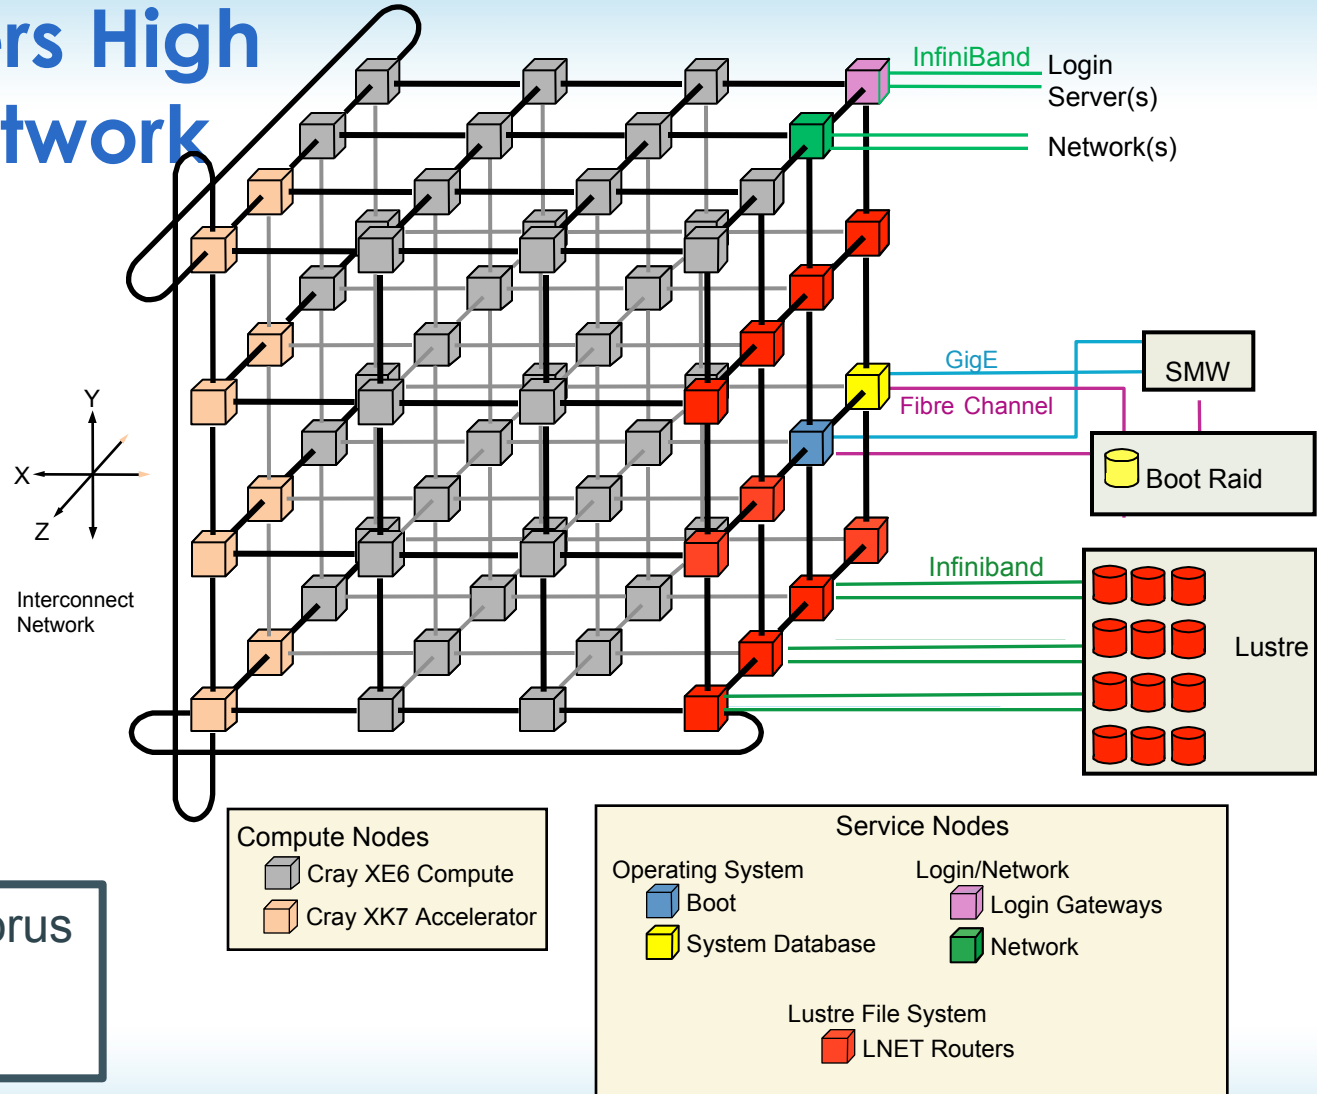
- Cray Sonexion with Lustre for all filesystems.
- All visible from compute nodes.

GO with Globus Online

- GridFTP client development for IE and HPSS nodes.
- Enabled data striping with GridFTP.
- Managed file transfers.
- Command line interface.
- Globus connect for sites without GridFTP endpoints.



Blue Waters High Speed Network




Blue Waters 3D Torus
Size
23 x 24 x 24

Blue Waters XE6 Node

Blue Waters contains 22,640 XE6 compute nodes

Node Characteristics	
Number of Core Modules*	16
Peak Performance	313 Gflops/sec
Memory Size	64 GB per node
Memory Bandwidth (Peak)	102 GB/sec
Interconnect Injection Bandwidth (Peak)	9.6 GB/sec per direction




**Each core module includes 1 256-bit wide FP unit and 2 integer units. This is often advertised as 2 cores, leading to a 32 core node.*

Cray XK7 and a Path to the Future

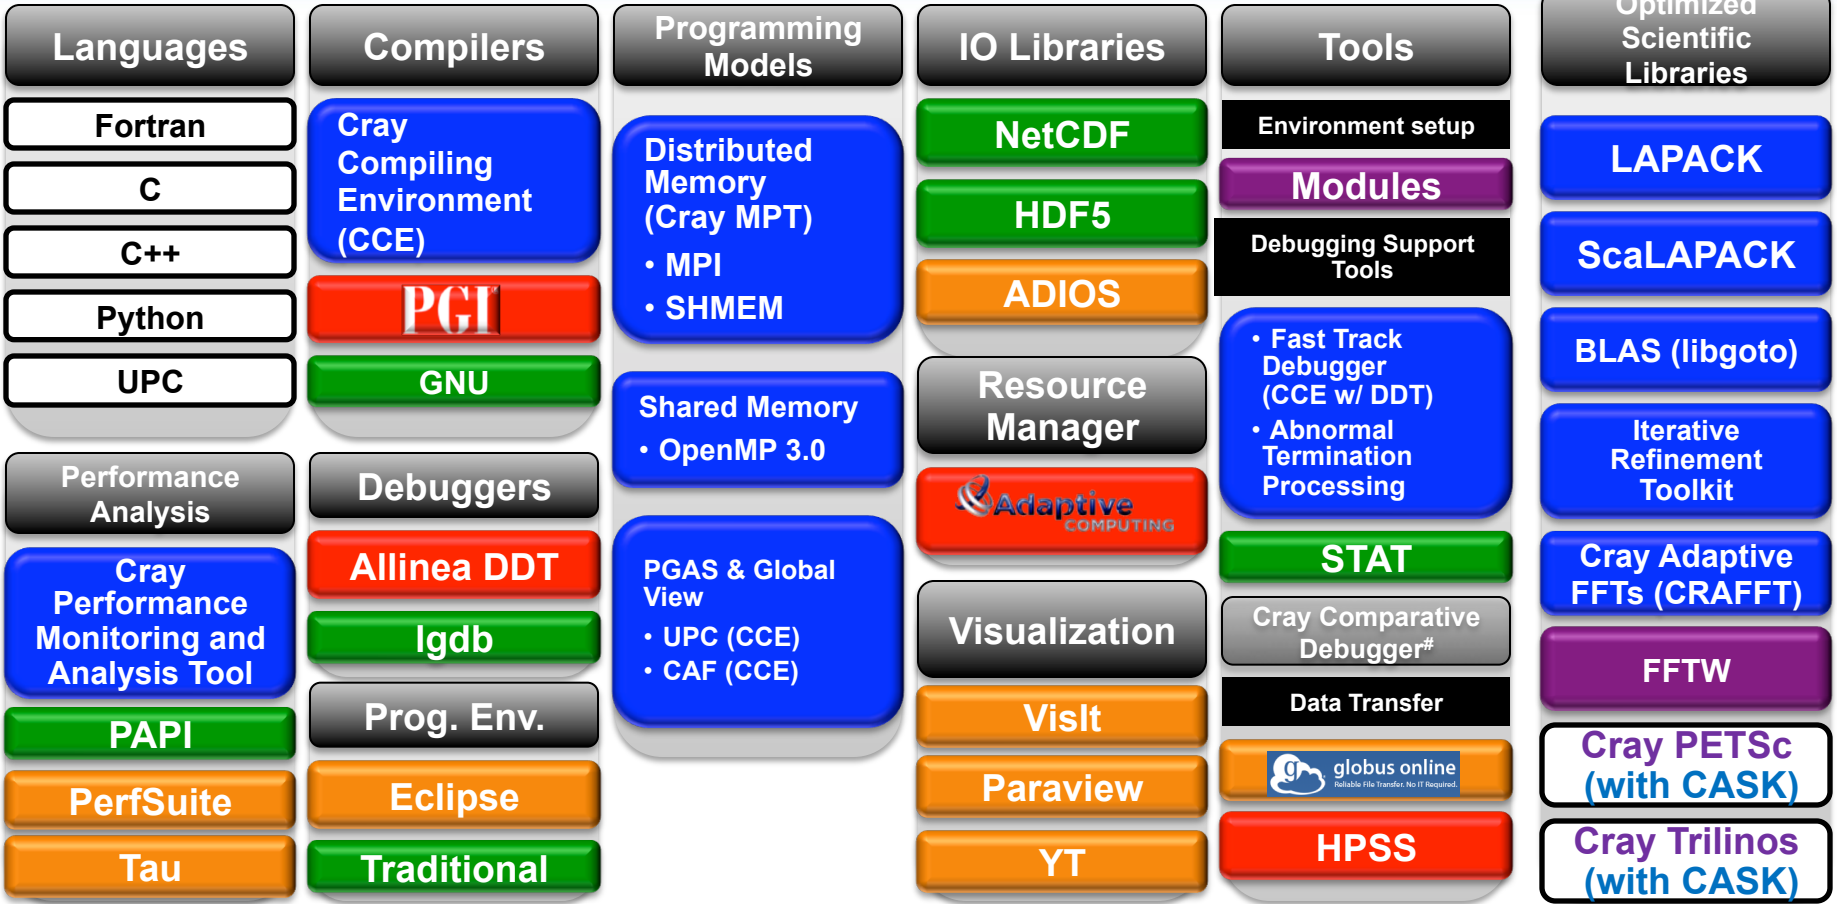
Blue Waters contains 3,072 NVIDIA Kepler (GK110) GPUs

XK7 Compute Node Characteristics

Host Processor	AMD Series 6200 (Interlagos)
Host Processor Performance	156.8 Gflops
Kepler Peak (DP floating point)	> 1.3 Tflops
Host Memory	32GB 51 GB/sec
Kepler Memory	6GB GDDR5 capacity > 180 GB/sec



Blue Waters Software Environment



Cray Linux Environment (CLE)/SUSE Linux

Cray developed	3rd party packaging
Under development	NCSA supported
Licensed ISV SW	Cray added value to 3 rd party

Summary

- Outstanding Computing System
 - The largest installation of Cray's most advanced technology
 - Extreme-scale Lustre file system with advances in reliability/maintainability
 - Extreme-scale archive with advanced RAIT capability
- Most balanced system in the open community
 - Blue Waters is capable of addressing science problems that are memory, storage, compute, or network intensive or any combination.
 - Use of innovative technologies provides a path to future systems
- Illinois/NCSA is a leader in developing and deploying these technologies as well as contributing to community efforts.