

BLUE WATERS

SUSTAINED PETASCALE COMPUTING

FujiFilm Global IT Summit

Michelle Butler

Senior Technical Manager

Storage And Integrated Network Technology (SAINT)

National Center for Supercomputing Applications, University of Illinois



GREAT LAKES CONSORTIUM
FOR PETASCALE COMPUTATION

CRAY

Who is NCSA

- NSF funded institution to research and provide cycles to USA academic researchers
 - Proposals submitted and reviewed
 - Cycles/storage granted
 - NCSA provides those cycles with consulting with storage and nearline needs throughout the life of the proposal need. (usually 1-2 years)
- Other sites funded: SDSC, TACC, NICS
 - Each site is unique in it's offerings




What do I do?

- I'm a saint. ☺
 - (Storage and Integrated Network Technologies)
- I lead all storage projects on all platforms. That is the research, design, implementation, and production of all HPC storage projects.
 - That includes the hardware and software for file systems, disk drives, RAID configurations, tape drives, libraries and library management systems.
- I lead also all networking projects on all platforms(with the help of a technical lead). That is the same as above, but for all networking platforms. I'm new to this (> 1.5 yrs)
- Actually I have a great technical people that lead themselves; self starters, researchers, and I get obstacles out of their way and enable them to continue



What is the need for SuperComputing ?

- Most USA researchers have closet clusters. Handfuls of machines stuffed into a closet .. Maybe a rack of 20-30 systems.
- As applications and the bounds of the science problem grows – so does the research size of the problem that is being analyzed.
- Department resources->campus resources->small NSF clusters->supercomputers->BlueWaters



BLUE WATERS

SUSTAINED PETASCALE COMPUTING



GREAT LAKES CONSORTIUM
FOR PETASCALE COMPUTATION

CRAY

NCSA

Private Sector Partners



industry.ncsa.illinois.edu

Machines at NCSA currently

- Iforge: Industrial partner machine
 - Partners machine (previous slide)
 - GPFS ½ PB file systems
 - 149 nodes
 - AMD cores – dual and quad socket
- GPFS condo
 - 1.5PB of shared storage
- Uofl campus cluster
 - 527 blade nodes
 - Intel (dual socket 8 core)
 - GPFS ½ PB file systems
- BlueWaters



Blue Waters System

- A NSF proposal in 2007-2008 Track 1 – 200 Million for machine alone
- Original partner IBM -> Cray
 - Within 3 months of contract had hardware on site
- 60 Million on building with 4-5 M retro fit for Cray
- 20M on networking and nearline storage environment
- Meant for LARGEST projects/users. Favors large jobs while small jobs sit and wait


BLUE WATERS

SUSTAINED PETASCALE COMPUTING



GREAT LAKES CONSORTIUM
FOR PETASCALE COMPUTATION

CRAY



Cray System & Storage cabinets: • >300

Compute nodes: • >25,000

Usable Storage Bandwidth: • >1 TB/s

System Memory: • >1.5 Petabytes

Memory per core module: • 4 GB

Gemini Interconnect Topology: • 3D Torus

Usable Storage: • >25 Petabytes

Peak performance: • >11.5 Petaflops

Number of AMD processors: • >49,000

Number of AMD x86 core module: • >380,000

Number of NVIDIA GPUs: • >3,000



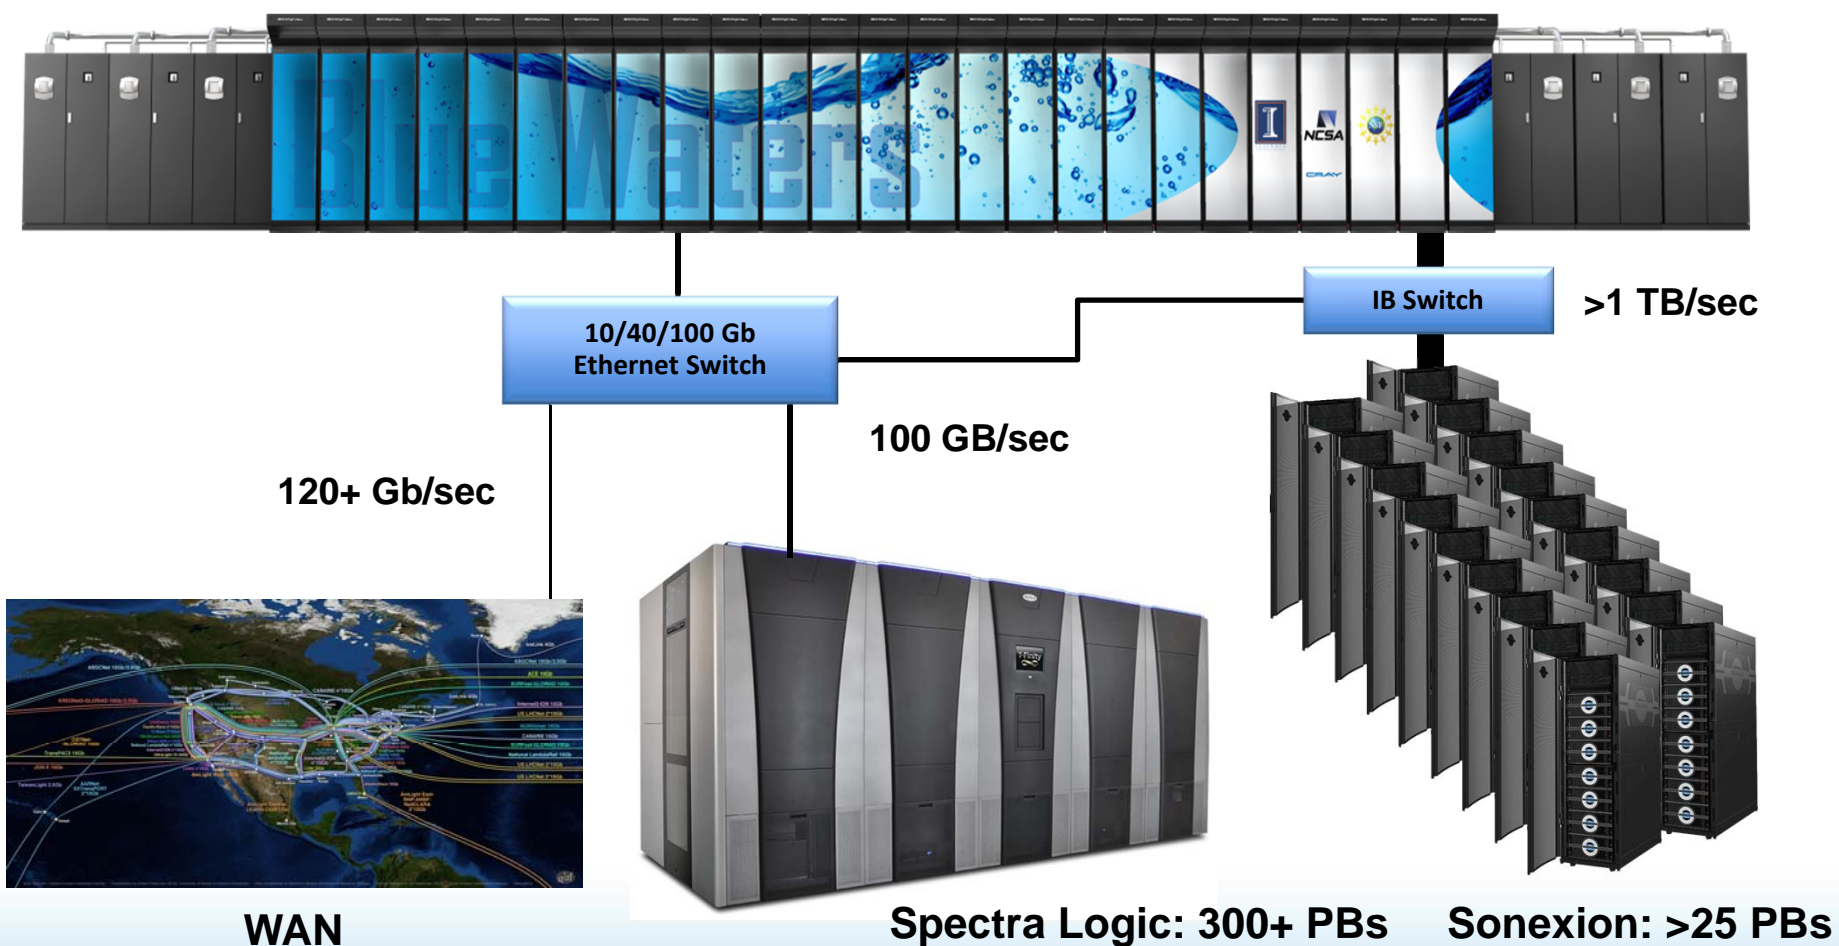
ILLINOIS
UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN

National Petascale Computing Facility




- Cooling towers
 - DC power
 - 24MW into building
 - Liquid cooled BW
-
- Modern Data Center
 - 90,000+ ft² total
 - 30,000 ft² raised floor
 - 20,000 ft² machine room gallery
 - Energy Efficiency
 - LEED certified Gold
 - Power Utilization Efficiency, PUE = 1.1–1.2

Blue Waters Computing Super-system




1 Minute Build






**More than 25 PRAC science teams
12 distinct research fields
selected to run on the new Blue Waters
Expect ~10 more major teams**




Nanotechnology



Astronomy/Astrophysics




Earthquakes and the damage they cause



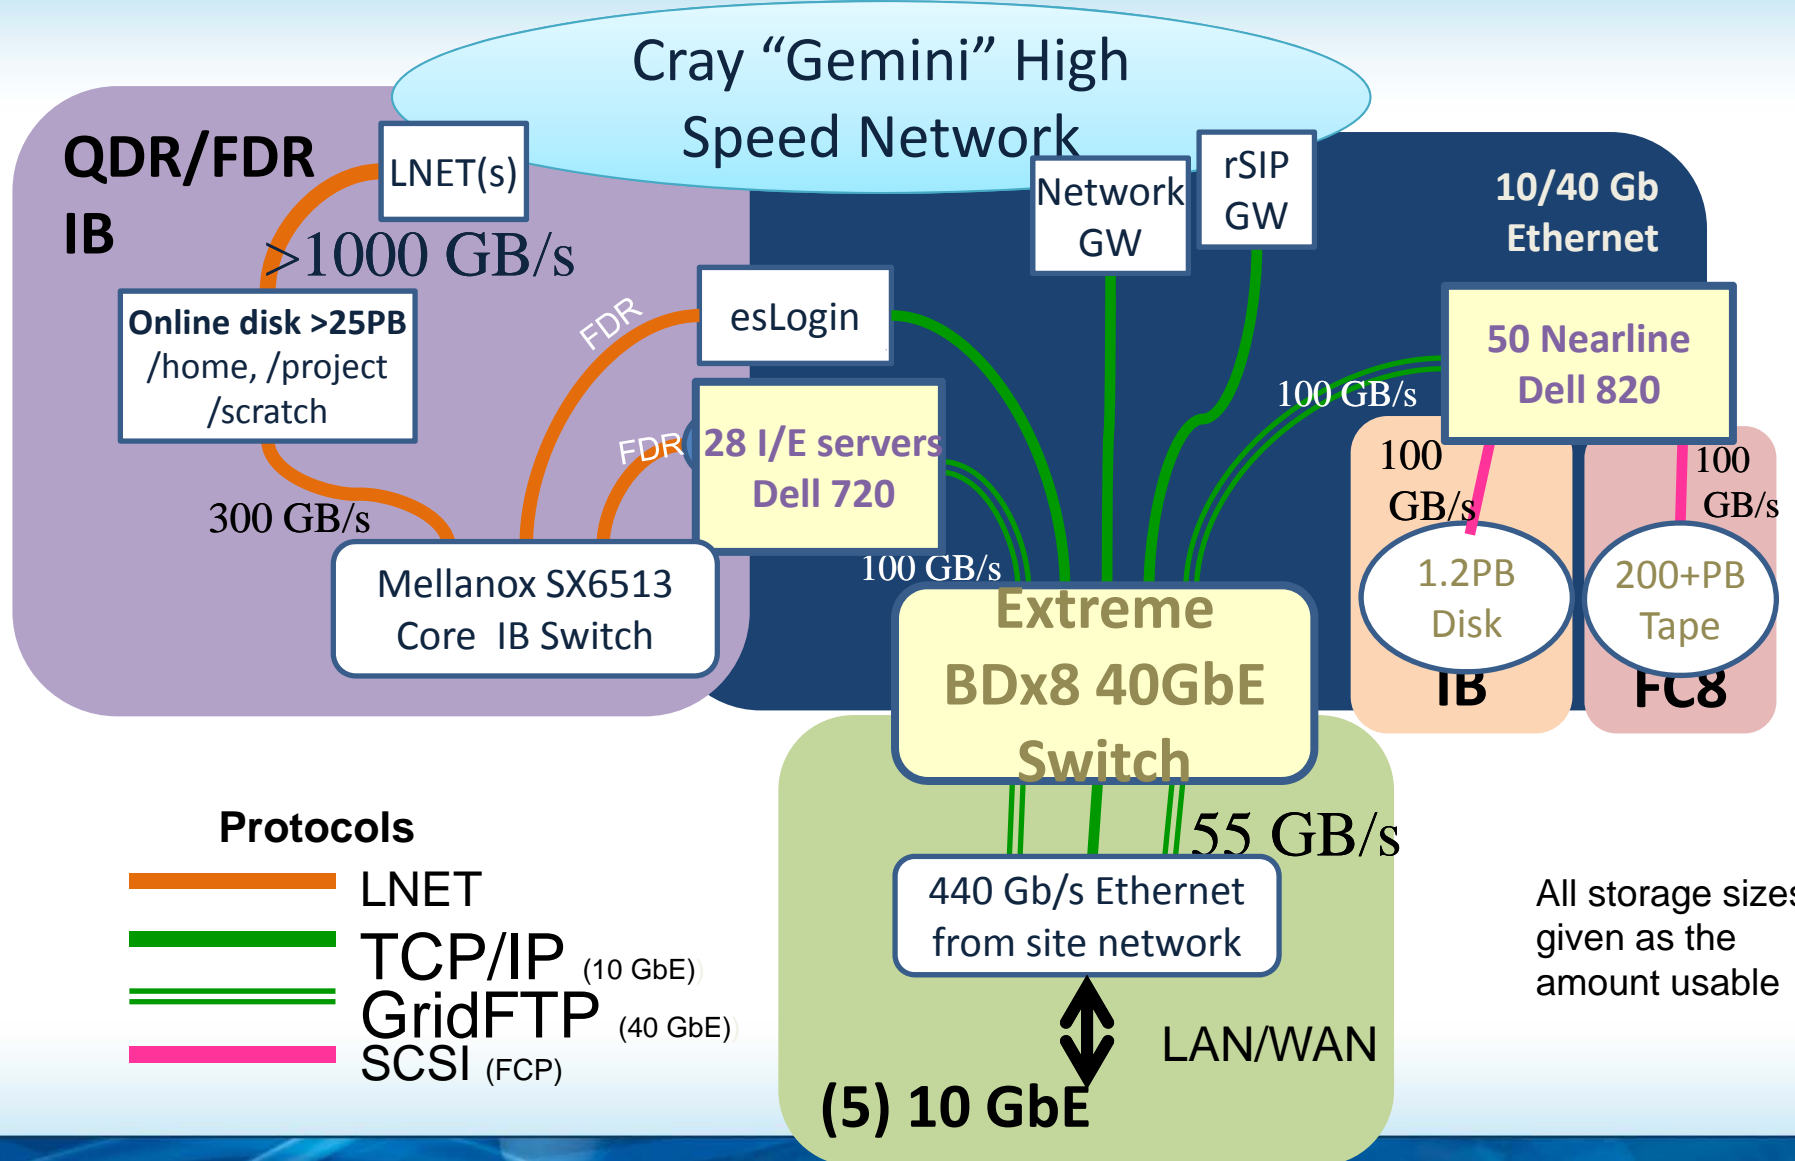
Viruses entering cells




Severe storms



Climate change




High Resolution Atmospheric Component of CESM: Projected Changes in Tropical Cyclones



High resolution (0.25°) atmosphere simulations produce an excellent global hurricane climatology

Courtesy of
Michael Wehner, LBNL




First Unprecedented Result – Computational Microscope



- Klaus Schulten (PI) and the NAMD group - Code NAMD/Charm++
- Completed the highest resolution study of the mechanism of HIV cellular infection.
- May 30, 2013 Cover of *Nature*
- Orders of magnitude increase in number of atoms – resolution at about 1 angstrom


Petascale Simulation of Turbulent Stellar Hydrodynamics

- Paul Woodward PI – Code PPM
 - 1.5 Pflop/s sustained on Blue Waters
 - 10,5603 grid
 - A Trillion Cell, Multifluid CFD Simulation
 - 21,962 XE nodes; 702,784 interger cores; 1331 I/Os; 11 MW
 - All message passing and all I/O overlapped w. comput.
 - 12% theoretical peak performance sustained 41 hrs
 - 1.02 PB data written and archived; 16.5 TB per dump.
 - Ran over 12 days in 6-hour increments



Enabling Breakthrough Kinetic Simulations of the Magnetosphere via Petascale Computing

- Homa Karimabadi PI – Code PPM
 - Possible extreme solar storms could significantly disrupt many modern infrastructure systems
 - This project studies the initiation and transmission of the solar wind



Summary

- Challenges are
 - Monitoring the BW environment
 - Hardware of the cray, IE and HPSS nodes, all tape drives, libraries, nodes, disks, controllers, and switches
 - Predictive analysis on what is going to fail
 - Providing utmost resource for our users;
 - Making them successful makes us successful.
 - Easing the file transfer mechanism;
 - 100Gbit to NCSA building
 - What is the next funding opportunity?