

BLUE WATERS - GEO

MAPPING AND MODELING THE WORLD

Blue Waters Operational Data

Brett Bode, Assistant Project Director
brett@illinois.edu

 **ILLINOIS**
NCSA | National Center for
Supercomputing Applications



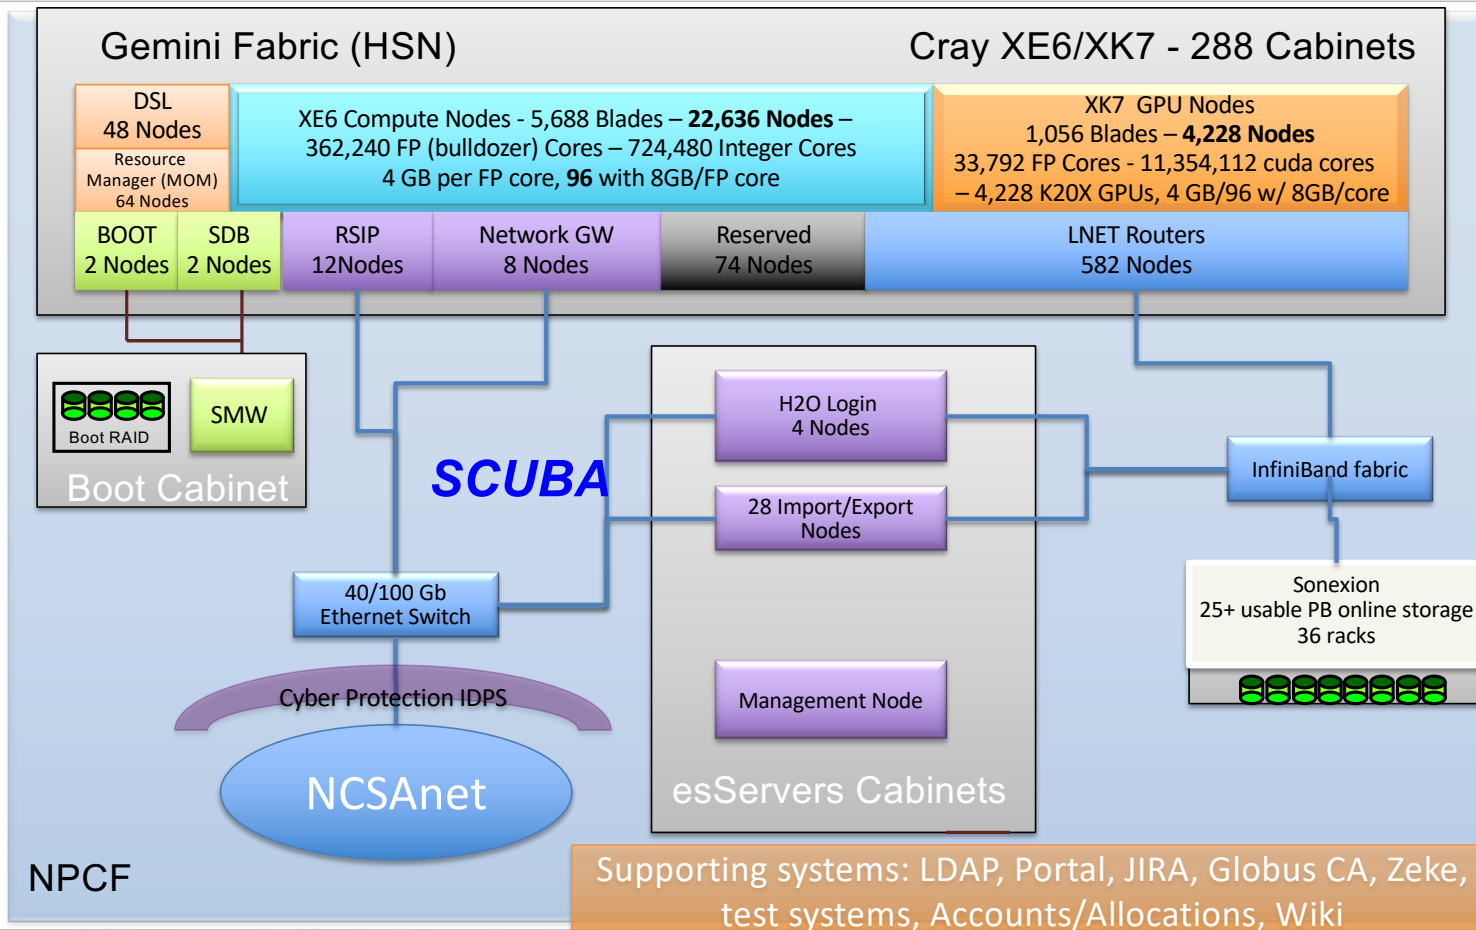
CRAY

UNCLASSIFIED

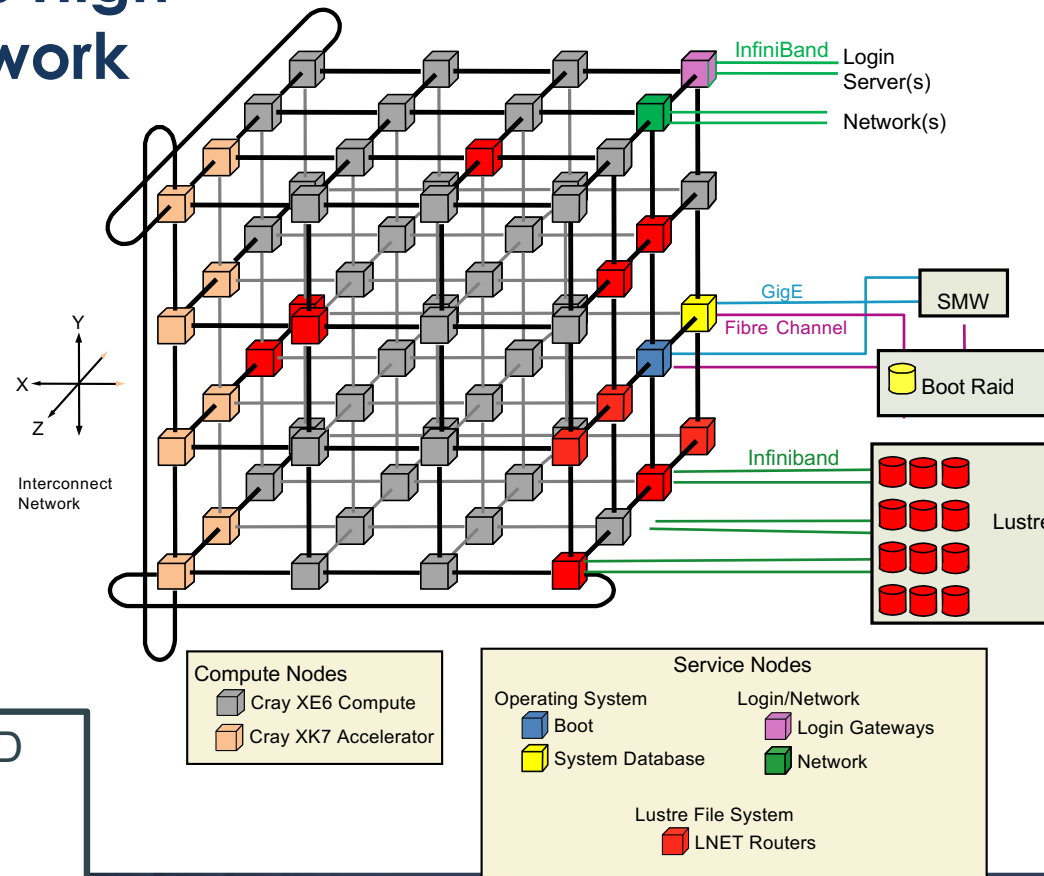
The Blue Waters Project

- Comprehensive development, deployment and service phases with co-design and other aspects
- Blue Waters is a **Universal Research Instrument** for frontier computational and data science that has delivered over 29.6 Billion core hour equivalents
- Make teams highly productive creating insights that cannot be done in other ways
- Diverse Science teams are able to make excellent use of those capabilities due to the system's flexibility and emphasis on sustained performance.
- The Blue Waters system is a top ranked system in all aspects of its capabilities. – **Intentionally not listed on Top500 list**
 - 45% larger than any system Cray has ever built. The next closest Cray system was #1 on the Top500 (Titan at ORNL).
 - Performance and delivered cycles are approximately the same as the aggregate of all the NSF XSEDE resources.
 - Ranks in the top systems in the world in peak performance – despite being over seven years old
 - Largest memory capacity (1.66 PetaBytes) of any HPC system in the world between 2012 and 2019!
 - One of the fastest file systems (>1 TB/s) in the world between 2012 and 2019!
 - Fastest external network capability (450 Gb/s) of any open science site
 - 147.4 Tb/s computer room network capability





Blue Waters High Speed Network




Blue Waters 3D
Torus Size
24 x 24 x 24

Lifetime Availability

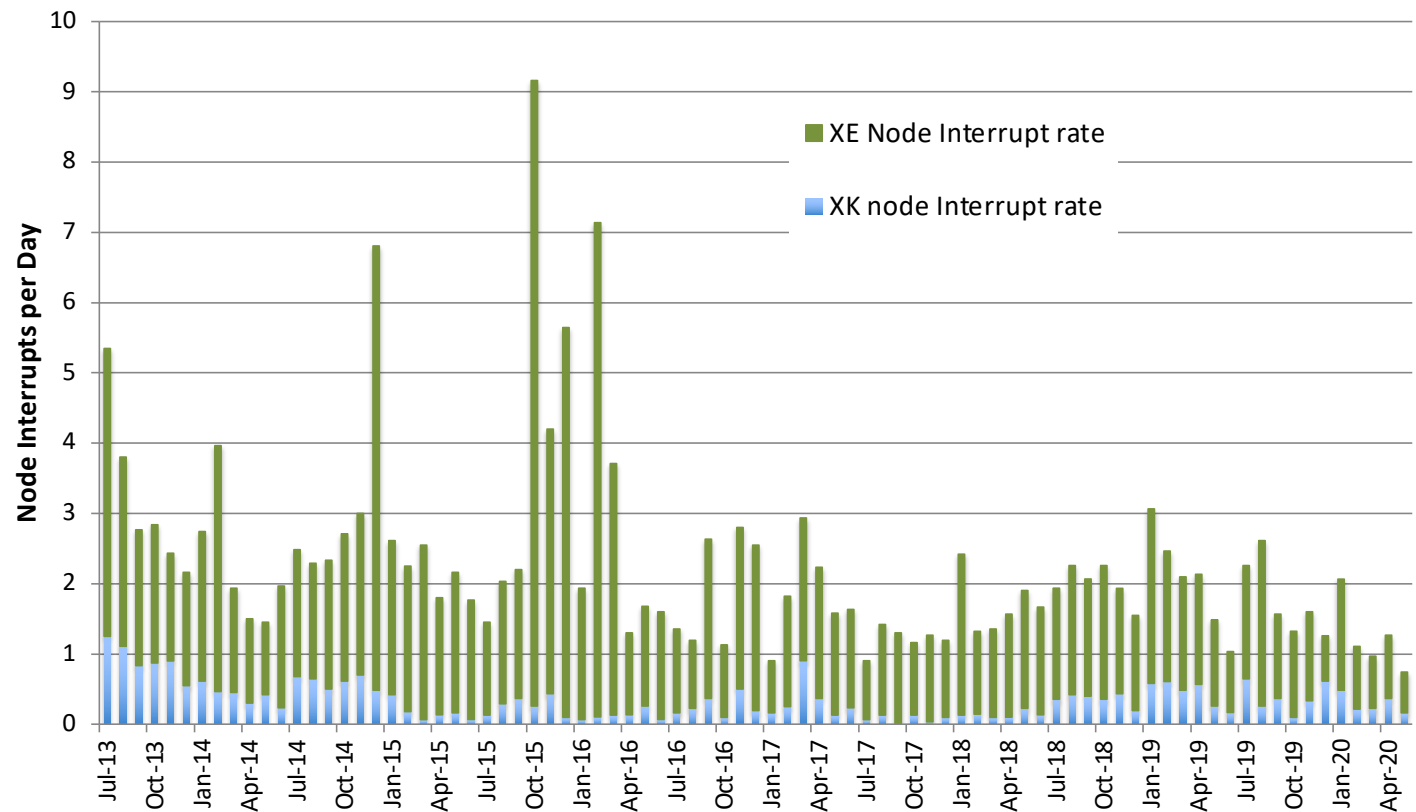
- With one exception BWs met its target availability throughout its service life.
 - Exception was period of subpar file system performance following an upgrade. Operation continued, but system was considered unavailable.

Scheduled Availability




Node Interrupt Rate

- The rate includes all sources of job interrupts.
- The highest peaks were caused by software issues.



System Wide Interrupts

- System wide interrupts are most often software related.
- System upgrades addressed multiple issues in the first two years.
- Averaging ~6 system outages per year the past three years.



Blue Waters Lustre File Systems

- Even after ~8 years the scratch sustained IO rate of 1TB/s is among the fastest in the world.
 - 26PB total usable space provided by >17,000 2TB hard drives
 - Very good at block IO!
 - DEM processing is sustaining >300GB/s
 - Only okay at file system metadata operations
 - If your workload opens/closes many, many files constantly talk to us, there are higher-performance options (that also impacts other users less).
 - Not as fast as your laptop SSD on a single node, but scales quickly and sustains the performance across the full system.




Hard Drives Fail Constantly

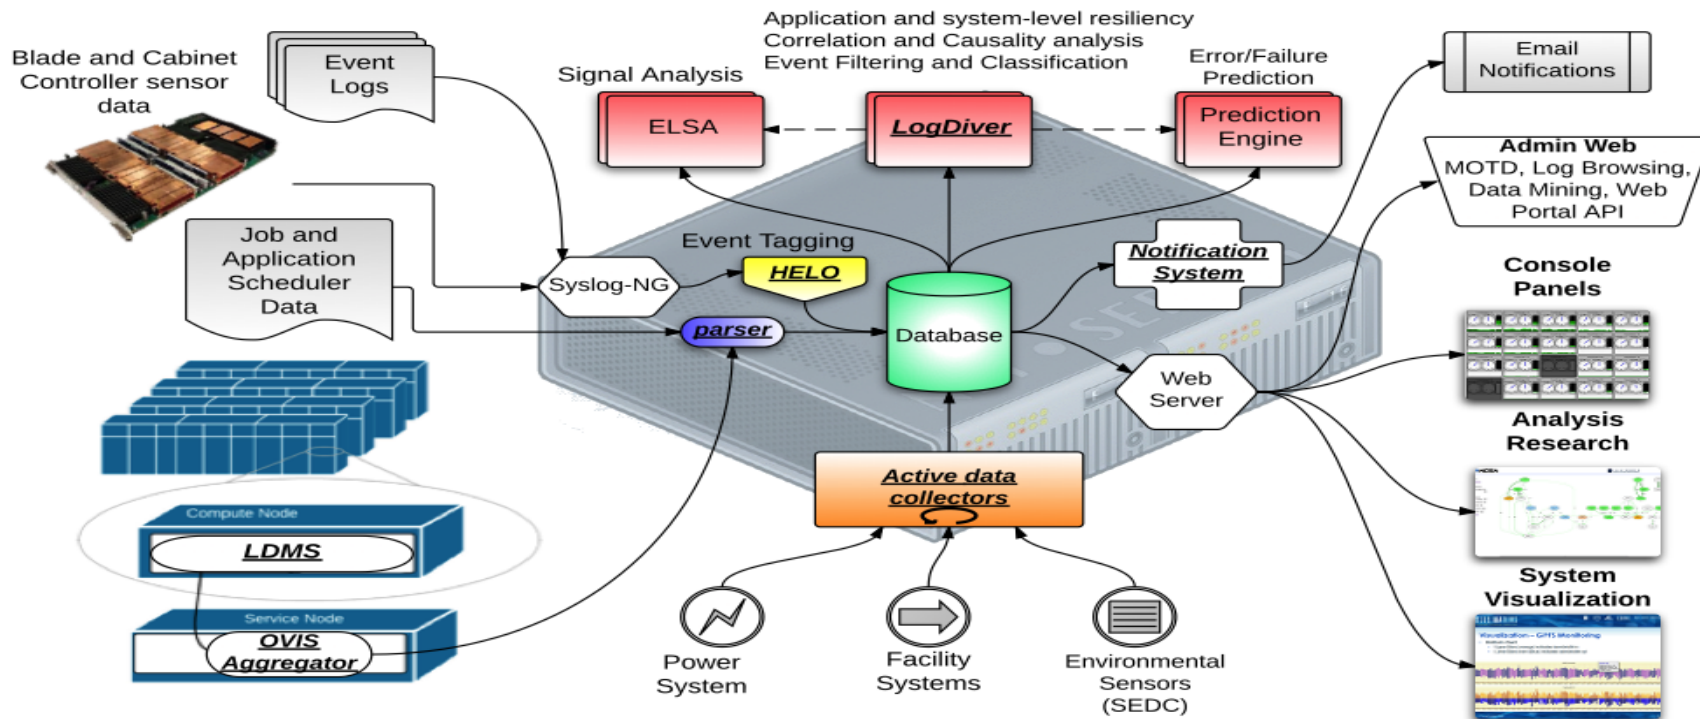
- The Blue Waters file systems were installed in Early to mid 2012.
 - ~ eight years of continuous operation!
 - Some drives fail in a catastrophic manner
 - Others exhibit many soft, recoverable/correctable errors.
 - The rate of these is often rapidly increasing.
 - Some drives exhibit many high-wait times for data return.
 - Both modes cause performance issues!

Number of 2TB Drive Failures and Preemptive Removals per Month

■ Failures ■ Preempt



Integrated System Console (ISC)



Data descriptions

- File stores
 - About 90G of MSR data, 55G of node metric data per day
 - Data is raw format such as counters
- Integrated System Console Database
 - ~16G per day of node data per day, 4 days retained
 - Some data is preprocessed from counters to rates
 - Example: flop counters per core are converted to flop rate per node
 - Greatly improves query efficiency
 - All log and system data
- Access methods
 - SQL queries
 - CSV files in lustre (parallel tools to extract data)
 - Web interfaces (with image/raw data downloads)
 - Published the datasets
- System Impact
 - Collection transfer and storage a 1 second intervals validated at full scale on Blue Waters without application performance impact
 - See publications page for reference to "[Lightweight Distributed Metric Service: A Scalable Infrastructure for Continuous Monitoring of Large Scale Computing Systems and Applications](#)"

Current Data

- Gemini Link Statistics
 - All 6 directions
 - Link BW, %used, average packet size, %input queue stalls, %credit stalls, ...
- Gemini/NIC Statistics
 - totaloutput_optA/B, total input, FMA output, bet output
 - SMSG
 - Number tx/rx rate , Bytes tx/rx rate
 - RDMA
 - Number tx/rx rate , Bytes tx/rx rate
 - IP over Gemini
 - Transmit/Receive rate
- Application library use (XALT)
- MPI I/O operations from each application (Darshan)
- Node
 - Load average
 - Latest,5min, running processes, total processes
 - Flop rate
 - Current free memory
 - GPU
 - Utilization, memory used, temperature
 - Pstate, Power Limit, Power Usage
- Filesystems
 - For each home, projects, and scratch
 - Bytes/sec Read and write
 - Rate of Opens, closes, seeks

Blue Waters Data is Available to Researchers!

- Blue Waters data is available for download via Globus.
- <https://bluewaters.ncsa.illinois.edu/data-sets>
- What is available:
 - Collected metrics
 - System Logs
 - Torque accounting logs
 - Darshan job data
 - Lustre user experience I/O data
- Username, id, and project codes are anonymized.

Questions?

- Email brett@Illinois.edu