

BLUE WATERS

SUSTAINED PETASCALE COMPUTING

Testing at Scale

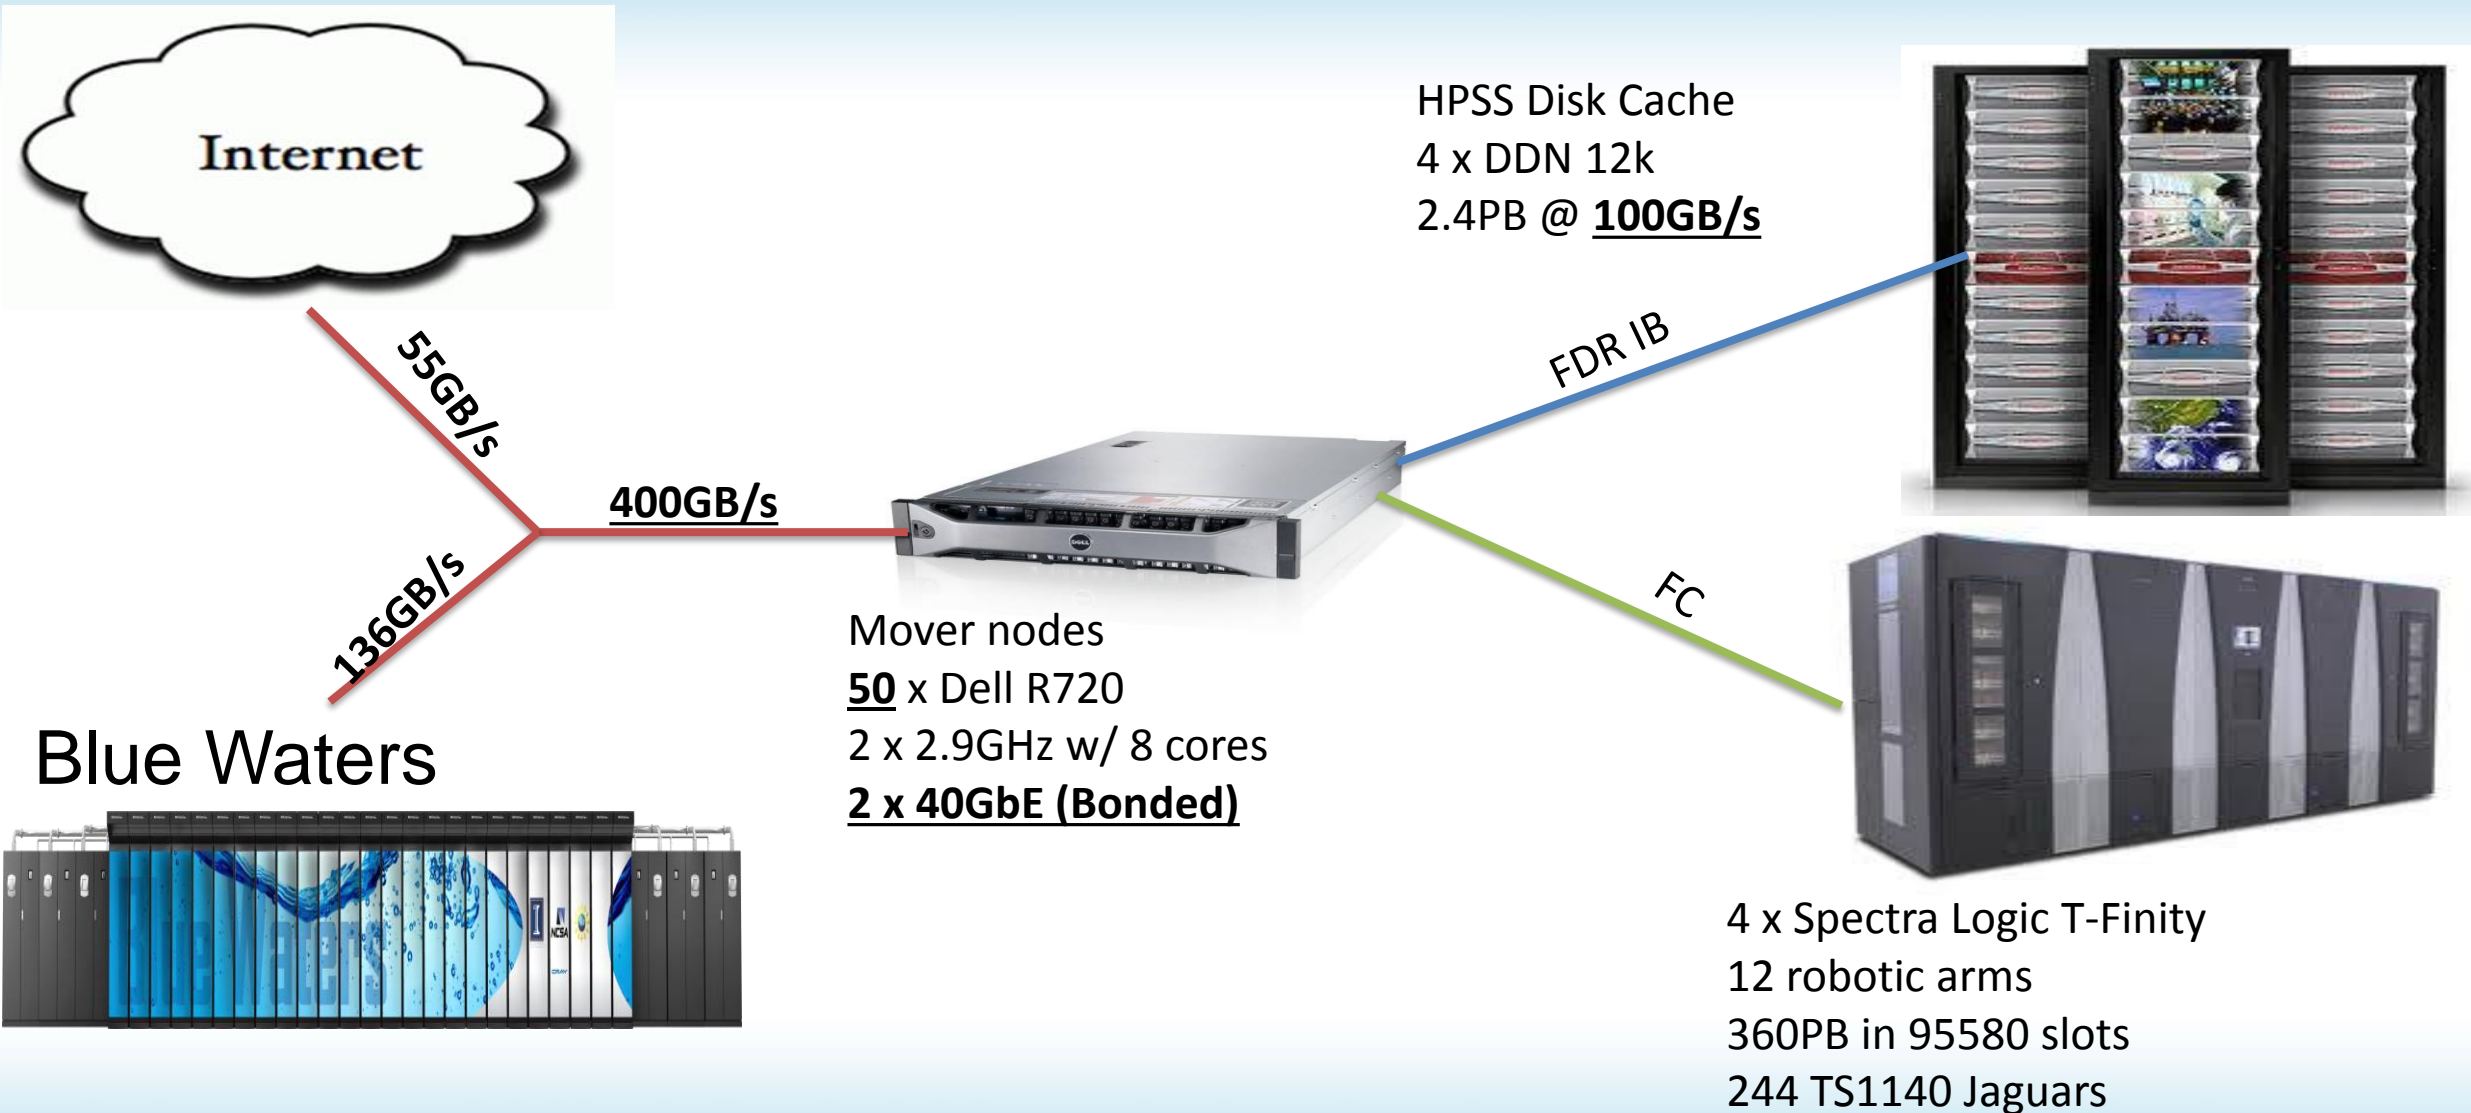


GREAT LAKES CONSORTIUM
FOR PETASCALE COMPUTATION



UIUC/NCSA AND CRAY
CONFIDENTIAL

Do not copy or distribute without expressed permission
from the NCSA Blue Waters Project Office



Develop System Use Cases (know what to test)

- Sustained throughput
- File creation and deletion rates
- Primary access mechanism
- Account allocations
- High availability / Failover

**Derive requirements. I want 100GB/s
throughput (who doesn't)**

Develop Ancillary Use Cases (performance or functional)

- Security considerations
- System logging
- Disaster recovery
- Service monitoring

Create an exhaustive list. Missed items will cost you later!!!

Break It All Down

- Prioritize based upon number of dependencies
 - Don't start with account allocation or primary interface testing
- Start with independent layers and build up
 - Each disk drive before virtual disk before LVM before file system
- Divide requirements by hardware quantity
 - $100\text{GB/s} / (\#\text{disks} * (D / (D+P))) = \text{necessary individual disk rate (X)}$
- Aggregate components to meet requirements
 - Each disk enclosure capable of $100\text{GB/s} / \# \text{ enclosures}$

Example Testing

- 40GbE baseline testing with iperf
- IB testing with native tools (ib_ping, ib_rdma_lat, ib_rdma_bw)
- Single drive, virtual disk, LVM testing with dd
- File system testing with XDD
- Tape drive performance with dd
- Add in DB2 testing, tape mounting, HPSS installation
- Primary interface testing (GridFTP)
- Wrap up with multi day continuous stress testing with all layers

This should all be obvious ... Your admins can do this is their sleep. Something important is missing and it's YOUR responsibility!!!

An Example of Failed Testing to Illustrate

- Buy several Dell 720s, 40GbE. Run iperf.
 - Stable 39Gb/s. Perfect. Next!
- Buy several Dell 840s, bonded 2x40GbE. Run iperf.
 - Variable 30-40Gb/s rates. Wait. What?
 - Try a new OS image. Nope. Try new drivers. Nope.
 - Contact RedHat. /shrugs /points_finger Contact Mellanox, Dell, etc
 - Exhaust brainpower. Done. Let timeline slip. Done. Panic!
- A valuable tool for successful large system deployment is missing



For use with cutting-edge deployments!!

Advice for System Designers

- Automate testing - it will be repeated
- Keep the design homogeneous – less to test
- Insist on vendor-provided testing – they are the experts
- Limit the number of (support) vendors – limits finger pointing
- Get an 800lb gorilla – for all those stubborn technical issues

Thank You!