# Applications driven System Architecture: SUN's HPC strategy

**Dr. Giri Chukkapalli, Ph.D**
**Sun Microsystems, Inc.**

# What will be your company's HPC success in 2012?

- Architect, design well balanced, petascale systems that are energy efficient and adhere to *Open standards and Opensource*
- Innovate within the Open Standards and OpenSource
- Do not equate Open Standards to Off the shelf
- Open standards therefore not a capability system

# Open Standards, Opensource

- X64?, Infiniband, Open  Storage
- Opensource
  - Studio compilers
  - ClusterTools
  - Lustre
  - GridEngine
  - Integrated HPC stack
  - Clean OS
- Broader incorporation of commercial technologies into HPC space: hardware, software and system management (cloud)

# Sun Magnum Supercomputer/HPC Data Center

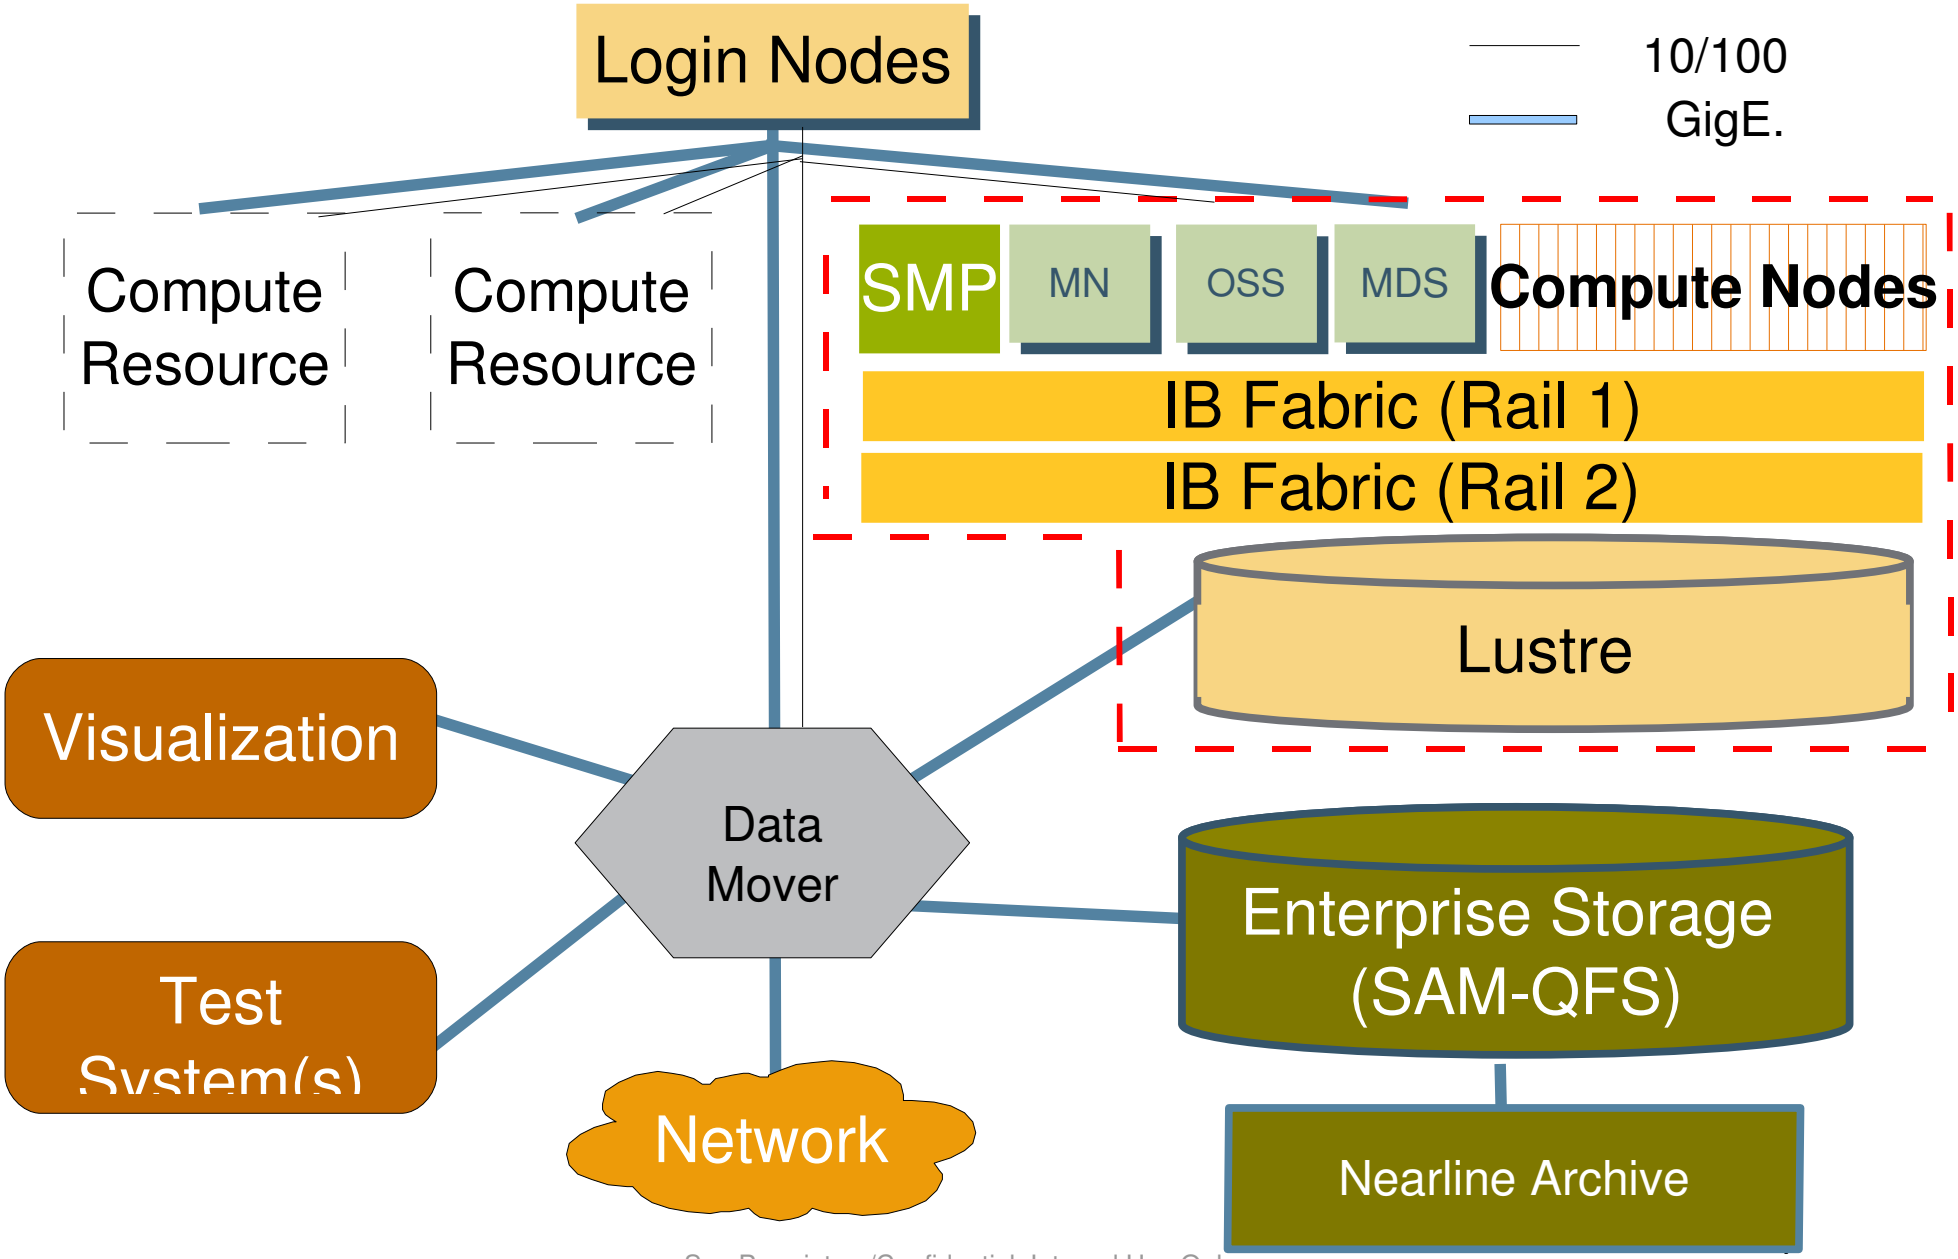

# TACC Ranger status

- 579TFLOPS
- Much better than expected MTBF (except USB drives)
- Extremely stable given the system size
- Atleast 8 projects that run at 32k cores
- Lot of science success stories from NSF, NOAA etc.

# 100PF soon after 2012?

- About 2 cores per year
- In addition FLOPS/cycle double every 3 years giving ~doubling every year
- 2PF in 2009 to 128PF in 2015
  - General purpose systems
- Aggregate performance
  - Increase Instruction Level Parallelism
  - Threads
  - Cores
  - Sockets
  - nodes

# Spectral parallelism

- Coupled multi-phase multi-physics application
  - Parallelism across physics components
  - Data parallelism in each physics component
    - synchronous
  - Functional parallelism
    - asynchronous
- Future supercomputer
  - Parallelism across asymmetric fabric
  - Inter-nodes DM parallelism
  - Intra-node SMP parallelism
  - SIMD parallelism
  - Instruction level parallelism

# HPC Industry challenges

- Customers will be deploying systems that cannot be powered in its full capability
- Idle power reduction
- Elevate memory and Interconnect to first class citizens in the system
  - Break free from pincount barrier
  - Move data between CPU, memory and Interconnect
  - Use concept of data/metadata or QoS

- Memory and CPU RAS
  - All the low hanging RAS already builtin
  - Hardware deals with metadata RAS application deal with data RAS
- Applications have to tag the data with QoS requirements
- HPC community is behind the curve in the concept of QoS
  - RAS
  - Data movement
  - Criticality of parts of code

- The result of economic meltdown
  - Adaptive system designs with organic efficiencies
  - If the system application efficiency is 5% of peak, it will use no more than 10% of power
  - Fine grained dynamic power management
- HPC application and system software is behind the curve in terms of asynchronous programming and resilient programming
  - Google solves similarly tightly coupled problems
  - Erlang code can achieve nine nines resiliency

# What customers should be doing?

- Break free from synchronous programming
  - Many of the performance problems go away
  - Jitter issue becomes somewhat irrelevant
  - RAS can be handled cleanly
- There is enough memory, Interconnect and IO BW but not for burst synchronous needs

# What's SUN's  strategy to grow

- Well balanced portfolio up and down the hardware/software stack and technical, commercial domains
- Not build “one of” systems
- Non religious system architectures
- Exploring global addressing, GPGPUs and other special purpose processors