
Integrated and Optimized Dell HPC Solution Bundles

Ramesh Rajagopalan

Enterprise Product Group, Bangalore

ramesh_rajagopalan@dell.com

Oct, 2010



Agenda

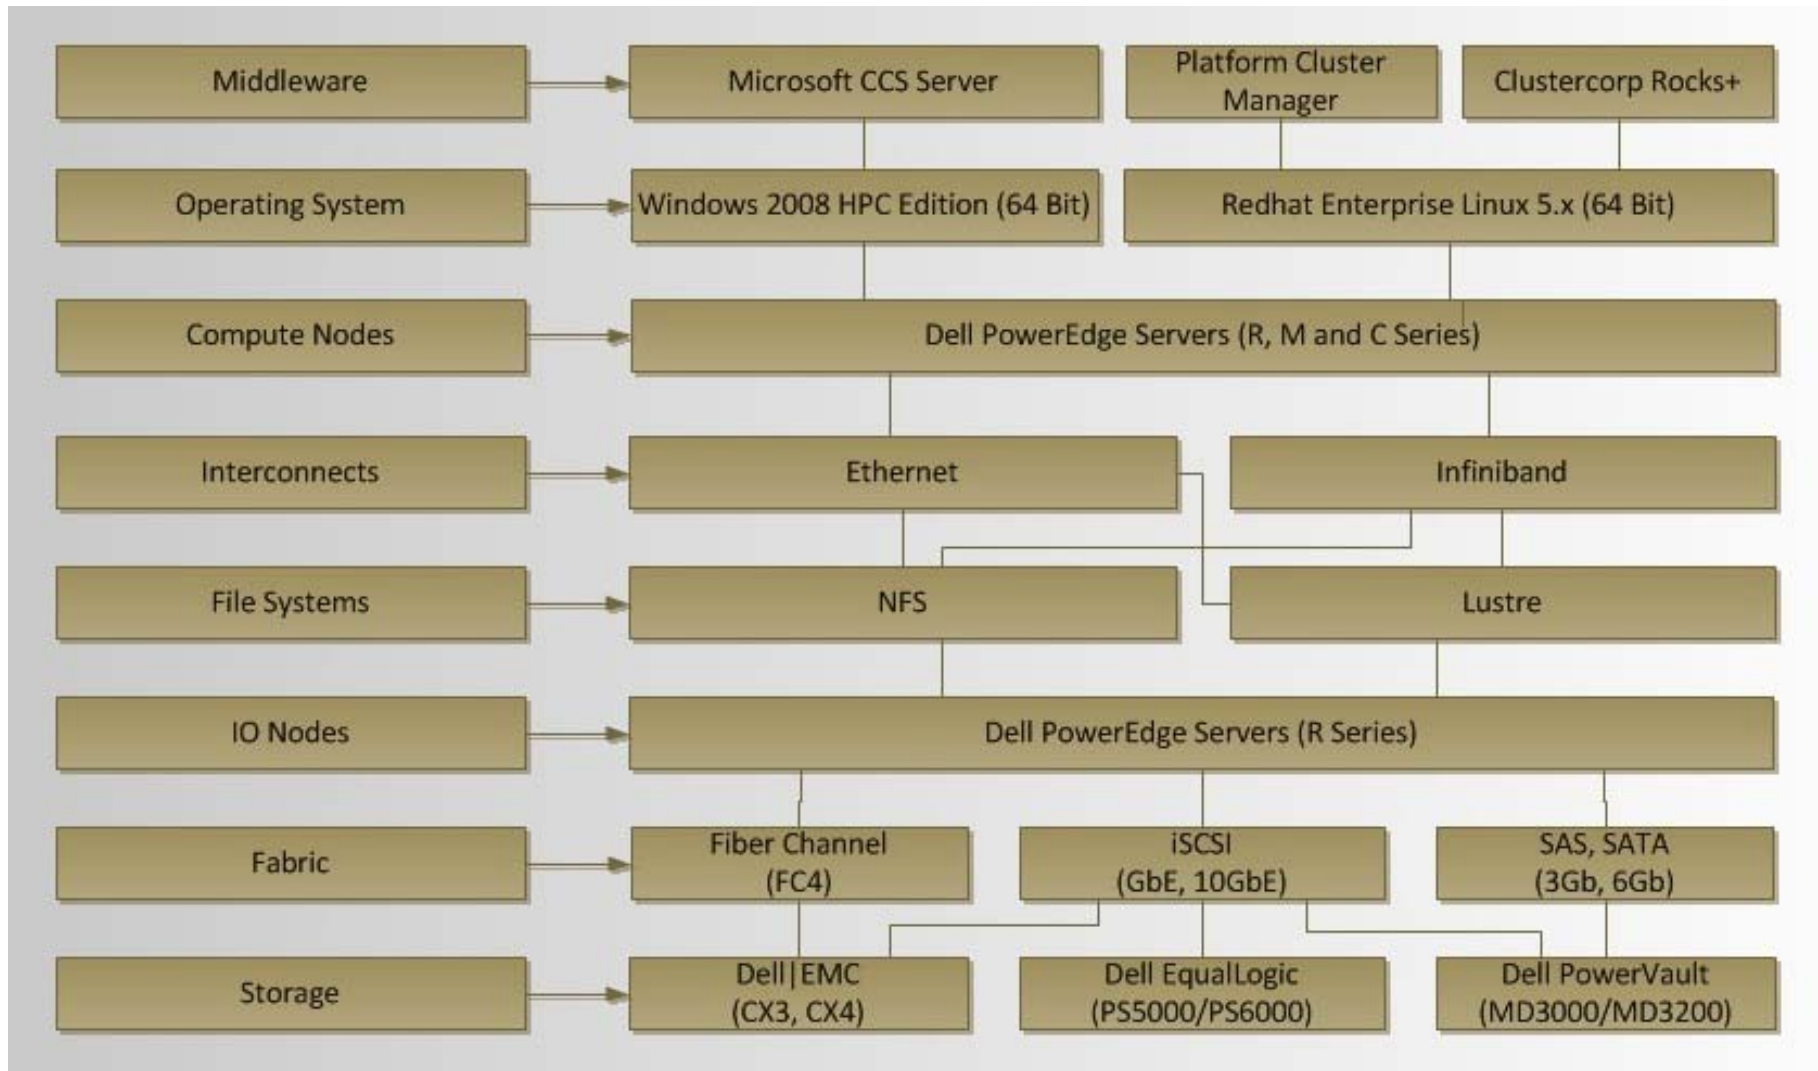
- Deployment
 - Server dependent
 - Diskless Compute Nodes
 - Optimized BIOS setting
 - Configuring Remote Management

- Advisor (Sizer Tool)

- Storage and File System



HPC Building Blocks



Deployment – Scenario # 1

- Different generation of Servers or Interconnects
 - dependent drivers or management packages
 - dependent methods to configure remote management
- Server specific drivers and server management packages are bundled
- During deployment specific driver or server management package is selected based on the kind of the Server
- Scripted tool to view and configure remote management or BIOS setting in 1- to-many fashion

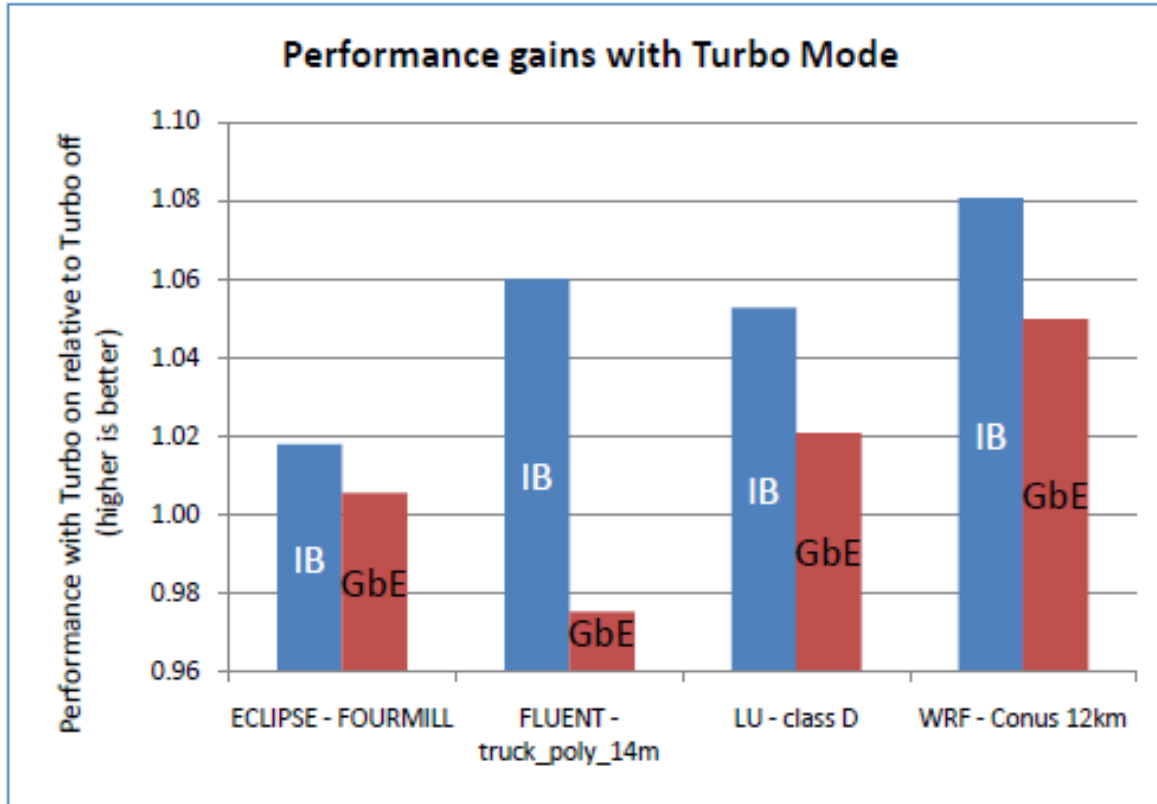


Deployment – Scenario #2

- **Different family of Servers**
 - optimized BIOS settings
- BIOS settings associated with Intel processors
 - C- state, P- state, Turbo mode
- BIOS settings associated with AMD processors
 - HT- assist, node- interleave, power profiles
- Subset of commonly used HPCC benchmarks and applications are run with various combinations of settings
- Common denominator of BIOS are identified and set as defaults



Deployment – Simplification #2 (contd)



Deployment – Scenario #3

- **Driver updates for diskless nodes**
 - Compute nodes are PXE booted and obtain the image from the master node
 - Driver update on the compute node will be lost after a reboot

- Scripted procedure to update the driver in the master node
 - Backs up the previous version for restore
 - Builds the new drivers
 - Updates the drivers and rebuilds the diskless image



Advisor

- Sizing runs with a set of HPCC benchmarks and applications
 - Different building blocks i.e., Servers, Interconnects, Storage
 - Different applications

- Sizing data and associated logic goes into the development of the Advisor Tool



Advisor Sample Sequence

HPC Advisor

For additional information or to order your customized solution, visit www.dell.com or call your sales representative.

COMPUTE CLUSTER

Number of Servers	128
Compute Node Server Model	PowerEdge R410
Front-end (Head node, Master Node) Server Model	PowerEdge R710

I/O CLUSTER

Number of I/O nodes	8
I/O Node Server Model	PowerEdge R710
File System Management Node Model	PowerEdge R710
I/O Fabric	Gigabit Ethernet

MPI FABRIC

MPI Interconnect	QDR Infiniband
Switches	Four 36-port QDR InfiniBand switches.
MPI Network Interface Card	Gen 2 Pci-E single-port memory-free 8X InfiniBand card

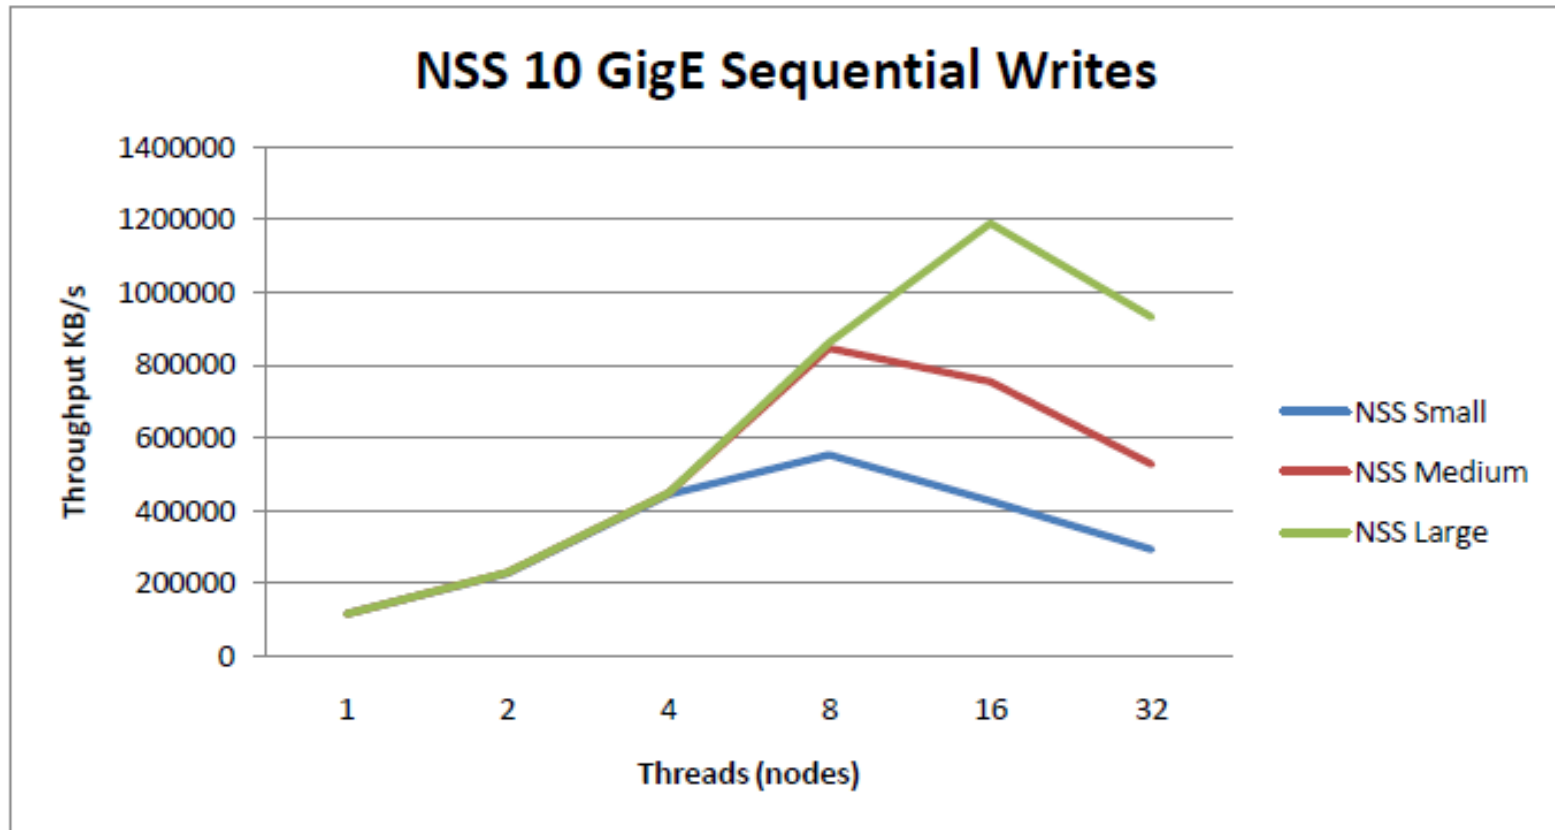
Storage and File system

- NFS and 10GigE
 - Small cluster i.e., 4 – 64 nodes
 - 10GigE or IB QDR yielding 1.1GB/s or 1.3 GB/s aggregate throughput
 - RedHat XFS
 - MD1200



<http://content.dell.com/us/en/enterprise/d/hpcc/storage-dell-nss.aspx>

Storage and File system – NSS Sample



References

- www.dell.com/hpcc
- http://content.dell.com/us/en/enterprise/d/business~solutions~whitepapers~en/Documents~HPC_Dell_11g_BIOS_Options.pdf.aspx

