

SGI® Solutions for SIMULIA Abaqus Customers

Tony DeVarco
Director of Manufacturing Solutions

Scott Shaw
Principal Applications Engineer



©2015 SGI


SGI Company Proprietary



Agenda

- SGI Solutions for CAE
- HPC Product Line Overview
 - SGI® UV™, SGI® ICE™ & SGI® Rackable®
 - SGI's Industry Partners and Customers
- Performance Tuning
 - Abaqus/Standard & Explicit
- SGI® VizServer®- Remote Visualization Solution
- Questions

Solutions for CAE-Compute, Viz & Storage



Typical CAE Workflow

- CAD model creation
- Mesh-generation
- Model decomposition
- Running solvers
- Viewing results
- Adjust and repeat

Data Management Options

1. Scratch Storage
 - High Performance-disk
2. Nearline Storage
 - Lower performance/ cost disk
3. Persistent Storage
 - Virtualized Storage Tiers

SGI® UV™ 2000

- Up to 64 sockets (512 cores), 16 TB shared memory/rack
- Intel® Xeon® Processor E5-4627 v2 8-core (3.30 GHz)
- NVIDIA® Tesla® and Intel® Phi™ support
- NVIDIA® Quadro K5200 or K6000 graphics support
- SGI NUMALink™ 6 Interconnect
- SUSE or RedHat
- Altair PBS Professional Batch Scheduler
- SGI® VizServer™ available for remote visualization



SGI® UV™ 2000

Individual Rack Unit (IRU) and Blades



Notes

- IRU: 10U high by 19" wide by 27" deep
- 8 blades – 16 sockets – per IRU
- 3 UV1 12V Power Supplies
- Nine 12V cooling fans N+1



**Compute
Blade**
Front




Accelerator Blades

Reference Configuration SGI® ICE™ X

SGI ICE X Achieves Top 4 out of 5 Most Efficient Supercomputers on TOP500 List *(June 24, 2014)
SGI ICE X systems can achieve 20 to 40+% higher performance efficiency vs. IBM, Bull, HP and Cray systems.**

- IP-113 (Dakota) blade for “D-Rack”
- 172 teraflops per rack of 1728 processor cores
- Intel® Xeon® E5-2690 v3 Twelve Cores 2.6GHz
- Houses up to two 2.5” SATA drives for local swap/ scratch usage
- Memory per core guidance
 - Implicit CSM codes- 4-8GB per core; Explicit CSM codes- 2-4GB per core; CFD codes 2-4GB per core
 - Integrated Infiniband FDR interconnect Hypercube/Fat Tree
 - Single or Dual-plane network topology
 - Multi-rail network (MPI communications isolated from the NFS traffic, splitting of large messages across 2 rails)
 - Integrated shared storage available
 - SUSE or RHEL with SGI Performance Suite



SGI® ICE™ X

Reference Configuration SGI® Rackable™

- Rackable C2112-GP2 (head node);
SGI C2112-4GP3 (compute nodes)
 - Intel® Xeon® E5-2690 v3 Twelve Cores 2.6GHz (135W)
- Local drives:
 - Head node- 500 GB SATA mirror Raid 1
 - Nodes- two 500 GB SATA striped for local swap/scratch usage. We also support SAS and SSDs
- Memory per core guidance
 - Implicit CSM codes- 4-8GB per core; Explicit CSM codes- 2-4GB per core; CFD codes 2-4GB per core
 - Infiniband QDR or FDR interconnect
 - Integrated shared storage available
 - SUSE or RHEL with SGI Performance Suite



SGI® Rackable™
20u and 42u
Racks Available

***Integrated and fully tested in the SGI
factory. Not on your Data Center!***

SGI's Partnerships

- **Strong 20+ years Industry Partnerships**
 - Continuing to work together in solving customer challenges
 - Working together in tuning Abaqus FEA on SGI Platforms
- **Established and Growing Industry Presence**
 - Large install base with Industry leading customers in HPC
 - Worked on many joint projects including HPC Experiment, remote visualization, user group conferences & benchmarking



Partial list of SGI's CAE Customers



Honeywell

NORTHROP GRUMMAN



Hamilton Sundstrand
A United Technologies Company



Sikorsky



Gulfstream®
A GENERAL DYNAMICS COMPANY



GENERAL ATOMICS



McLaren
AUTOMOTIVE

HONDA



EXEDY
EXEDY Corporation

TOYO TIRES®



Valeo



SUMITOMO



MAGNA

DENSO



BRIDGESTONE

3M



Performance Tuning

SGI Performance Suite

- MPI/OpenMP Communication Layers:
 - Optimization of Hybrid MPI and OpenMP threads
- Abaqus Process Thread Placement:
 - Autodetection of UV SMP and distributed clusters to bind application threads to mitigate migration across the system
- Improve batch Job efficiency:
 - Create batch scheduler topology resources so jobs run more efficient on various SGI platforms
- Improve Abaqus/Standard execution environment:
 - Leveraging SGI tools to reduce turnaround times by using SGI Performance Suite and SGI's Perfboost MPI layer intercept for both SGI UV SMP distributed clusters
- Abaqus/Explicit execution environment:
 - Was enhanced to use the SGI Performance Suite with SGI's Perfboost MPI layer

Why SGI Perfboost for Abaqus

- SGI PerfBoost and SGI MPI provides an enhanced MPI communication layer and tools for tight control of process placement in hybrid MPI/OpenMP environments
- The bundled IBM Platform MPI (default MPI) implementation for Abaqus FEA lacks the necessary tools for thread affinity when scaling to large core counts on clusters and SGI UV Systems.
- The PerfBoost shim library only intercepts and optimizes MPI calls while other application routines execute without intervention. The mathematical solution is the same but with faster turnaround times.
- Using SGI PerfBoost on SGI UV Systems can be 2x or more faster and about 15-32% faster on ICE X Clusters.

SGI Perfboost Results on ICE X Clusters

- The SGI Custom abaqus_v6.env site file is transparent to the user
- The user can still override or add settings to the Abaqus environment using a local abaqus_v6.env file
- Works with Abaqus 6.14 and below releases to Abaqus 6.8


Purple line is SGI tuning over IBM PMPI using -aff=automatic:bandwidth -affcycle=numa

Blue line is SGI tuning over the Abaqus Defaults

Automotive Model on ICE X Cluster


E5-2690 v3 2.60GHz 128GB Memory, Abaqus 6.13-3

1.7M Elements, 6.8M DOF, No. Steps 3, Iterations 52, FLOPs/Iter ~3.70E+12



SGI Perfboost Results on UV SMP Systems

- The SGI Custom abaqus_v6.env site file is transparent to the user
- The user can still override or add settings to the Abaqus environment using a local abaqus_v6.env file
- Works with Abaqus 6.14 and below releases to Abaqus 6.8



SGI Perfboost Improves Abaqus Scaling

Processor Choices & Runtime Comparisons

Intel® E5-2600v3 Series Processor


Intel E5 Series Processor Comparison with Abaqus v6.14-1

Processor comparisons

- E5-2667 v2 8c @ 3.20GHz 20MB Cache
- E5-2680 v3 12c @ 2.50GHz 30MB Cache
- E5-2690 v3 12c @ 2.60GHz 15MB Cache


New Abaqus Models

- S2b – Flywheel with Centrifugal load 1.8M DOF FLOP/iter 4.78E+13
- S8 – Shipping port heavy contact 3.2M DOF FLOP/iter 8.82E+12
- S10 – Plate with uniform load 9.9M DOF FLOP/iter 1.08E+13
- S11 – Block with uniform load 9.9M DOF FLOP/iter 4.82E+13



SGI VizServer for Remote Visualization

- Secure Web Portal that allows users to interact with backend architectures
- Web Portal provides system resource monitoring of backend resources, batch scheduler, NVIDIA GPUs and Intel Phi accelerators
- Round-Robin device scheduling through the batch scheduler



Problem:


- The HPC users were expecting their result data from their new SGI UV to cause them to spend up to 4 – 8 hours in data transfer back to their workstations for analysis.
- They were looking for a way to streamline their workflow.

Solution:

- 4 NVidia® Quadro K2000 graphics cards in the UV with VizServer, coupled with custom scheduling integration to maximize the number of possible users at any given period of time.
- Allows for fast access to result data and post processing.

SGI VizServer Client Success

Aerospace Company




Back of the rack view

Abaqus on SGI® Tech Guide

Available for Download

<http://www.sgi.com/pdfs/4395.pdf>



Conclusions

- SGI Performance Suite can reduce runtimes by detecting the hardware environment correctly
- SGI Perfboost integration is transparent to Abaqus users and can be applied as a site environment file
- Take advantage of the new Intel E5 v3 Series processors which provides more cores per processor, memory bandwidth, AVX2 turbo instructions and Intel® Turbo Boost v2.0 features
 - E5-2680 v3 12c @ 2.50GHz 30MB Cache
 - E5-2690 v3 12c @ 2.60GHz 15MB Cache
 - E5-2667 v3 8c @ 3.20GHz 20MB Cache
- SGI Remote Visualization is an excellent way to cut costs at the desktop by accessing latest GPUs centrally on a UV 2000 Server

Questions