WHITE PAPER

SGI® ICE™ 8400 Powered by AMD Opteron™ 6200 Series Processors

Performance, Cost/Performance, and Scalability

TABLE OF CONTENTS

1.0 Introduction 3

2.0 SGI ICE 8400 with AMD Opteron™ Series Processors 3
- 2.1 Architected for Extreme Processor Efficiency 4
- 2.2 A Choice of Blades with AMD Opteron™ 6200 Series Processors 5
- 2.3 A Choice of QDR InfiniBand Interconnect Topologies and Chassis 5
- 2.4 Hierarchical System Management 6

3.0 AMD 2nd Generation Architecture and AMD Opteron™ 6200 Series Processors 7
- 3.1 Architectural Overview 7
- 3.2 AMD Turbo CORE Technology 8
- 3.3 Flex FP 256-bit FPU 8
- 3.4 New Technical Computing Instructions 9
- 3.5 Power Efficiency 9
- 3.6 Straight-Through Computing 10

4.0 Standard and Application Benchmarks 10
- 4.1 Standard Single-Node Benchmark Comparison 11
- 4.2 SPECfp_rate_base2006 and SPECint_rate_base2006 11
- 4.3 SPEC MPI2007 12
- 4.4 OpenFOAM 13
- 4.5 ANSYS Fluent 14

5.0 Conclusion 14

6.0 References 15

"The use of HPC and mining large datasets continue to grow at James Cook University (JCU). Our existing facilities were completely overloaded and we were running simulations the day the new facility was open to users that the old system could not handle."

— Ian Atkinson, Associate Professor at JCU.

1.0 Introduction

From energy exploration and manufacturing to large scientific simulations that push back the limits of knowledge, technical and high-performance computing (HPC) systems face familiar impediments. Daunting technical computing challenges and truly immense datasets exert a constant pressure on available computing resources. Even with steady advances in processors, memory, and systems, larger and larger distributed clusters are required. For some, addressing these infrastructure hurdles has resulted in sprawling and complex clusters that are difficult to deploy, expensive to manage, and often confront very real limitations in terms of physical data center real estate, power, and cooling.

In contrast, SGI ICE 8400 powered by AMD Opteron™ 6200 Series processors are designed for flexible and scalable clustered deployment, and can offer dramatic improvements in terms of density, performance, efficiency, and manageability. This paper describes the SGI ICE 8400 system as well as unique capabilities and advancements offered by AMD Opteron™ 6200 Series processors. Results from recent standard benchmarks and applications are also provided.

2.0 SGI ICE 8400 with AMD Opteron™ Series Processors

The SGI ICE 8400 blade system helps minimize overhead and communication bottlenecks that can rob efficiency and scalability, especially for data intensive work-flows. ICE 8400 combines the powerful AMD Opteron™ 6200 Series processor architecture with a unique board and interconnect design that together drive a host of compelling advantages.

- **Leading Density.** Industry-leading density helps reduce space, power, and heat, with up to 128 sockets and 2,048 processor cores hosted in a single data center rack.
- **Performance and scalability.** SGI ICE 8400 with AMD Opteron™ 6200 Series processors offer 33% more cores than AMD Opteron™ 6100 series in the same physical footprint[1], yielding 35% greater throughput and 33% more memory throughput[2]. Performance features for threaded workloads include straight-through computing to prevent bottlenecks and compromises.
- **Efficiency.** Efficiency is important for processing, power and cost reduction, and these systems offer 35% greater throughput within the same power budget of previous-generation systems2. C6 power gating and TDP power gap help provide more control over power while AMD Turbo CORE technology ramps frequency predictably to help meet computing needs.
- **Flexible storage choices.** SGI ICE 8400 provides flexible storage choices, including storage over InfiniBand, local storage on service nodes, and optional local blade storage.
- **Interconnect flexibility.** To meet application needs and provide cluster scalability, integrated switching and a choice of interconnect topologies is provided. Single or dual-plane QDR InfiniBand is supported, and clusters can scale from 32 to 65,536 nodes (1024 racks).

1 Comparison of 12-core AMD Opteron™ 6100 Series procesors and 12-core and AMD Opteron™ 6200 Series processors. SVR-5

2 Testing based on STREAM benchmark. 4P: 146GB/s using 4 x AMD Opteron™ processor Model 6276 in "Drachma" reference design kit, 64GB (16 x 4GB DDR3-1600) memory, SuSE Linux® Enterprise Server 64-bit 110GB/s using 4 x AMD Opteron® processors Model 6176 in "Drachma" reference design kit, 64GB (16 x 4GB } DDR3-1333) memory, SuSE Linux® Enterprise Server 11 64-bit. 2P: 73GB/s using 2 x AMD Opteron(tm) 6200 Series processors in "Dinar" reference design kit, 32GB (8 x 4GB DDR3-1600) memory, SuSE Linux® Enterprise Server 64-bit 55GB/s using 2 x AMD Opteron™ processors Model 6176 in "Dinar" reference design kit, 32GB (8 x 4GB DDR3-1333) memory, SuSE Linux® Enterprise Server 11 64-bit. According to AMD internal engineering estimates, the AMD Opteron™ 6200 Series processors will deliver up to 35% greater throughput in the same general power/thermal envelopes as AMD Opteron 6100 Series processors. SVR-26

- **Efficient cooling.** SGI ICE 8400 servers offer a choice of air or water cooling to provide flexibility or to integrate with site-specific systems.
- **Integrated management.** Hierarchical management framework controllers aggregate management data so that even very large clusters can be managed easily and efficiently.
- **SGI factory integration.** SGI offers factory integration for ICE 8400 servers, allowing organizations to deploy HPC clusters rapidly and with less possibility for error.
- **Visualization.** Optional GPU graphics and GPGPU accelerators can be easily integrated to fit specific visualization or computational needs.

2.1 Architected for Extreme Processor Efficiency

SGI ICE 8400 represents a no-compromise, high performance blade system for scale-out clusters. The system features a unique board and interconnect design that allows clusters to scale easily from 32 nodes to tens of thousands, to address the most challenging compute problems. Because the system is modular, it combines the advantages of low entry-level cost with scalability in terms of processors, memory, InfiniBand connectivity, and I/O.

ICE 8400 can be optimized for industry-leading performance or price/performance and features superior system configuration flexibility. Five different compute blades can be configured, giving organizations a choice of processor type, network topology, system fabric, and storage. All ICE 8400 systems feature tightly-integrated QDR InfiniBand interconnects, with options for both single-plane and dual-plane support. Both value- and performance-optimized configurations can be specified. The 42 rack unit (42U) rack shown in Figure 1 houses up to four individual rack units (IRUs), with each IRU housing up to sixteen individual single-wide compute blades. A minimal system consists of two IRUs.

Figure 1. Up to four ICE 8400 IRUs can be installed in a single 42U rack, offering up to 2048 cores with AMD Opteron™ 6200 Series processors.

The ICE 8400 platform is fundamentally architected to provide cost-effective high-performance InfiniBand based infrastructure. With AMD Opteron™ 6200 Series processors, the system is capable of achieving industry-leading scalability without sacrificing application performance efficiency. ICE 8400 offers a variety of interconnect options that let organizations scale their applications across hundreds or thousands of processor cores.

The SGI ICE 8400 system can accommodate up to 16 compute blades within each IRU. The IRU is a 10 rack unit (10U) chassis that provides power, cooling, system control, and network fabric for up to 16 blades via a backplane. Up to four IRUs are supported in each custom-designed 42U rack, with a choice of either air cooling or water cooling for all configurations. Each rack supports:

- A maximum of four IRUs
- Up to 64 compute blades (up to 128 AMD Opteron™ 6200 Series sockets, and 2048 processor cores)
- A maximum of 32.8 TB of memory (64 x 512 GB)
- A choice of compute blades
- Integrated QDR InfiniBand interconnects
- 5+1 redundant 1625W 12V DC output power supplies (6+1 when high-wattage sockets are configured)
- 7+1 redundant 175mm blowers

2.2 A Choice of Blades with AMD Opteron™ 6200 Series Processors

To accommodate a range of computational needs, SGI ICE 8400 offers a choice of IP-110 or IP-106 compute blades equipped with AMD Opteron™ 6200 Series processors. Both blades offer the following capabilities:

- Two sockets for AMD Opteron™ 6200 Series processors
- 16 DDR3 DIMM slots per blade
- 4, 8, 16, and 32GB 1600MT/s ECC registered DIMMs
- Support for a 2.5-inch SATA HDD or 1.8-inch SSD

Though the compute blades are slightly different physically, their basic architecture is nearly identical, differing only in the QDR InfiniBand connectivity options supported:

- IP-110 compute blades support one dual-port Mellanox ConnectX-2 InfiniBand HCA
- IP-106 compute blades support two single-port Mellanox ConnectX-2 InfiniBand HCAs

2.3 A Choice of QDR InfiniBand Interconnect Topologies and Chassis

InfiniBand fabrics present different advantages and limitations, and applications and organizations have different topology requirements. Rather than dictating topology choices, SGI offers a choice of InfiniBand topologies, along with a choice of chassis that add additional flexibility and cost effectiveness. Supported topologies include the following:

- **All-to-All.** All-to-All topologies are ideal for applications that are highly sensitive to Message Passing Interface (MPI) latency since they provide minimal latency in terms of hop-count. Though All-to-All topologies can provide non-blocking fabrics, and high bisection bandwidth, they are restricted to relatively small cluster deployments due to limited switch port counts.
- **Fat Tree.** Fat Tree or CLOS topologies are well suited for smaller node-count MPI jobs. Fat Tree topologies can provide non-blocking fabrics and consistent hop counts resulting in predictable latency for MPI jobs. At the same time, Fat Tree topologies do not scale linearly with cluster size. Cabling and switching become increasingly difficult and expensive as cluster size grows, with very large core switches required for larger clusters.

- **Standard Hypercube.** Standard Hypercube topologies are ideal for large node-count MPI jobs, provide rich bandwidth capabilities, and scale easily from small to extremely large clusters. Hypercubes add orthogonal dimensions of interconnect as they grow, and are easily optimized for both local and global communication within the cluster. Standard Hypercube topology provides the lightest weight fabric at the lowest relative cost with a single cable typically used for each dimensional link.
- **SGI Enhanced Hypercube.** Adding to the benefits of Standard Hypercube topologies, SGI Enhanced Hypercube topologies make use of additional available switch ports by adding redundant links at the lower dimensions of the hypercube to help improve the overall bandwidth of the interconnect.

SGI also offers a choice of ICE 8400 blade chassis designed to help optimize the performance and economy of particular topology choices.

- The 8400EX chassis provides higher performance and bandwidth per blade and reduces congestion for dual-rail MPI applications. Two or four 40 Gb/second QDR InfiniBand switch blades can be installed in a single 8400EX chassis, providing two high-performance planes.
- The 8400LX chassis represents an ideal price/performance solution where a single-plane network shares all types of traffic. One or two QDR InfiniBand switch blades can be installed in the 8400LX chassis offering one high-performance plane.

2.4 Hierarchical System Management

Unlike traditional flat clusters, the ICE 8400 system does not have a head node. Instead, the SGI ICE 8400 system of hierarchical management framework controllers help management scale in a seamless fashion with the cluster. Management data from individual IRUs and racks across the cluster is aggregated at a higher level, allowing administrators to drill down to acquire specifics. This approach provides for isolation of components as well as isolation of management and run-time functions, and allows for hot-swap components. The hierarchical management framework is organized in three tiers, including:

- **Tier1.** The System Administration Controller (SAC) provides the management for the cluster and fabric, and is used to install ICE system software, administer that software, and monitor information from all of the compute blades in the system.
- **Tier2:** The Rack Leader Controller (RLC) is generally used by administrators to provision and manage the system using SGI Cluster Management (CM) software. The RLC can also provide boot services and root file system services for compute blades if desired.
- **Tier3:** Chassis Management Controllers installed in each IRU provides IRU management and OS synchronization for the blades installed within the IRU.

Service nodes can also be added into the system to provide for login, batch, or gateway functionality. Larger multi-processor nodes can be added as required. Optional GPU and GPGPUs can be added for additional visualization or computational support.

3.0 AMD 2nd Generation Architecture and AMD Opteron™ 6200 Series Processors

With HPC clusters reaching the petaflop level today, price per gigaflop can have a drastic effect on the overall cost of systems and the size of cluster that can be achieved. Most organizations want to get the most computing power they can afford with their allocated budget. In addition, power efficiency is an increasingly essential factor for consideration. The amount of power required, along with the cost of that power, can make a significant difference in the long-term cost of HPC clusters. Reflecting a complete redesign, the AMD 2nd Generation architecture and AMD Opteron™ 6200 Series processor represent outstanding value in terms of both price per gigaflop and gigaflops per TDP Watt, offering both considerable performance and power efficiency and making them an ideal match for the strengths of the ICE 8400 system.

3.1 Architectural Overview

AMD Opteron™ 6200 Series processors provide the world's first 16 core x86 CPU, offering outstanding value and power efficiency in a modular architecture that is enabled for future innovation. The AMD Opteron™ 6200 Series processor offers key features, including:

- Flex FP – the world's only dynamic, flexible floating point complex
- A revamped integrated memory controller
- AMD Turbo CORE technology for additional performance
- Reduced power requirements and flexibility through TDP Power Cap

AMD's flexible approach provides for both shared and dedicated resources to help maximize power efficiency and costs. Dedicated components (shown in blue in Figure 2) help increase performance and scalability. Shared components (shown in green and gray) maximize performance per watt.

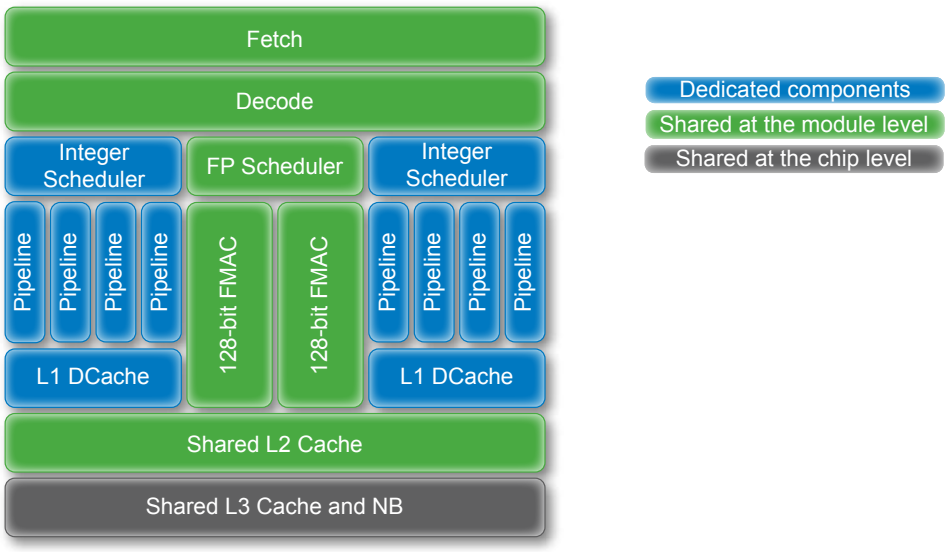

Figure 2. The AMD Opteron™ 6200 Series processor provides both flexible shared and dedicated resources to help maximize performance and power efficiency while minimizing costs.

3.2 AMD Turbo CORE Technology

Organizations need to get the most performance available from their systems, but they need that performance tailored to their application needs. AMD Turbo CORE technology provides multiple ways to increase the frequency of processor cores when there is sufficient "headroom" in terms of thermal design power (TDP), allowing AMD 6200 Series processors to deliver maximum gigaflops per TDP Watt. Figure 3 illustrates a single AMD Opteron™ 6200 Series processor in different states.

- **Base frequency.** At the left of the diagram, the sixteen-core processor is shown running at its base frequency, with available TDP headroom. In this state, there is essentially untapped potential frequency.
- **All Core Boost.** When there is TDP headroom in a given workload, AMD Turbo CORE technology is automatically activated, and can increase clock speeds by 300-500 MHz across all cores.
- **Max Turbo Boost.** When a lightly threaded workload sends half the modules into a C6 sleep state, but also requests maximum performance, AMD Turbo CORE technology can increase clock speeds by up to 1 GHz+ across half the cores in the chip.

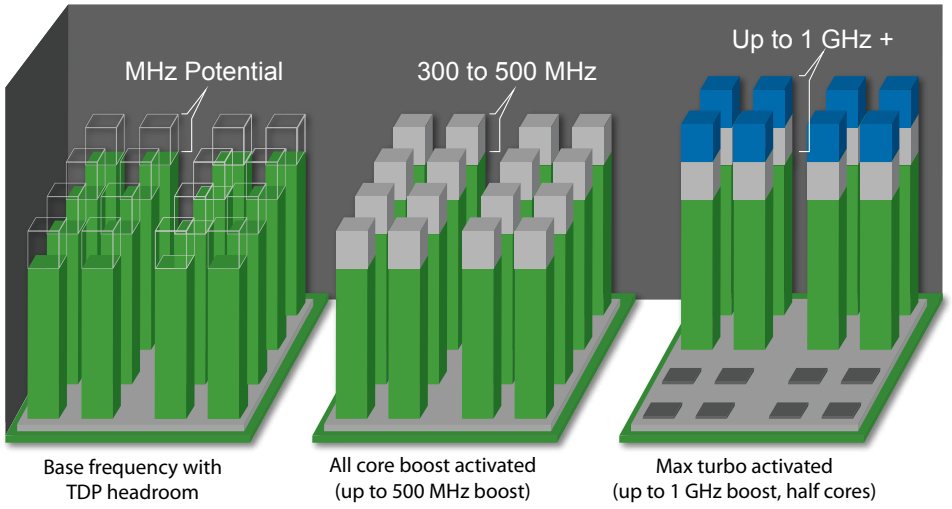

Figure 3. AMD Turbo CORE technology provides up to 1 GHz in additional processor frequency when there is sufficient TDP headroom.

3.3 Flex FP 256-bit FPU

Organizations need floating point performance, but it shouldn't come at the cost of excessive power consumption for applications that don't need it. AMD Opteron™ 6200 Series processors provide Flex FP, an innovative flexible floating point unit that is shared between two integer cores. Flex FP simultaneously executes two 128-bit instructions or a single 256-bit instruction and provides a dedicated floating point scheduler to help minimize latency for floating point applications. This innovative approach saves die space and conserves power for the majority of non floating point applications.

3.4 New Technical Computing Instructions

AMD Opteron™ 6200 Series processors provide support for new instructions that can help accelerate technical computing applications. Some of these instructions are also supported by other x86 chip vendors, while some are unique to AMD. Table 1 lists these new instructions along with applicable applications and use cases. Software that currently supports SSE3, SSE4.1, SSE4.2, AESNI, and AVX should simply run on systems equipped with AMD Opteron™ 6200 Series processors. Software will need to be written to access the AMD-specific instructions, or code will need to be compiled with a compiler that generates appropriate code utilizing these instructions.

Table 1. New technical computing instructions supported in AMD Opteron™ 6200 Series processors.

Instructions	AMD	Other x86 Vendors	Applications/Use Cases
SSE3, SSE4.1, SSE4.2	X	X	Video encoding and transcoding Biometrics algorithms Text-intensive applications
AESNI PCLMULQDQ	X	X	Applications using AES encryption Secure network transactions Disk encryption (MSFT BitLocker) Database encryption (Oracle) Cloud security
AVX	X	X	Floating point intensive applications Signal processing/seismic Multimedia Scientific simulations Financial analytics 3D modeling
FMA4	X	—	HPC applications
XOP	X	—	Numeric applications Multimedia applications Algorithms used for audio/radio

3.5 Power Efficiency

As clusters grow and organizations strive to get more out of their computing resources, efficiency can make the critical difference. With the completely new AMD 2nd Generation architecture, AMD Opteron™ 6200 Series processors provide up to 56% better power efficiency on a per-core basis than competitors[3].

- **More low power memory choices.** AMD Opteron™ 6200 Series processors offer a choice of low or ultra-low memory solutions.
- **Lower idle CPU power.** The C6 power state on AMD Opteron™ 6200 processors shuts down clocks and shuts off power to idle cores, reducing idle CPU power by 46%[4].
- **More power control flexibility.** TDP Power Cap enables more power control for IT, allowing the flexibility to set power limits without capping frequency.
- **Intelligent circuit design.** An all new intelligent circuit design helps minimize the number of active transistors for low power and outstanding performance

3 Based on AMD Opteron 4200 Series processor with 8 cores at 35W TDP versus lowest wattage, highest core Intel® Xeon® processor with 6 cores at 60W TDP according to www.intel.com as of November, 2011. SVR-79

4 Based on internal testing as of 8/2011: AMD Opteron™ processor model 6174 (12-core 2.2GHz) consumes 11.7W in active idle C1E power state, while AMD Opteron™ processor model 6276 (16-core 2.3GHz) consumes only 6.4W in the active idle C1E power state with new C6 power gating employed. System configuration: “Drachma” reference design kit, 32GB (8 x 4GB DDR3-1333) memory, 500GB SATA disk drive, Microsoft Windows Server 2008 x64 Enterprise Edition R2. SVR-60

3.6 Straight-Through Computing

HPC applications need massive amounts of throughput, both into and out of processors, as well as between processors, and between processors and memory. As deployed in SGI ICE 8400, AMD Opteron™ 6200 Series processors feature four full-speed HyperTransport connections between processor sockets (Figure 4). Each processor, in turn, has four full-speed memory channels to eight 1600 MHz DDR-3 DIMM slots. As a result, AMD can deliver "straight through" computing with dedicated pipelines for each integer thread, full speed memory on all models, and full speed between processors on all models.

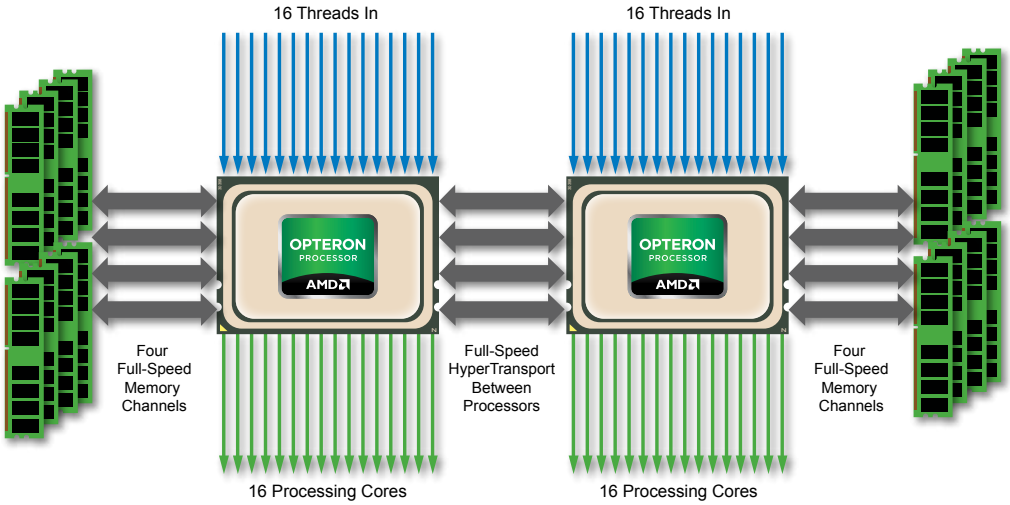

Figure 4. AMD Opteron™ 6200 Series processors provide straight-through computing, with 16 threads in and out of 16 processor cores and four full-speed links between processor, and processors and memory.

4.0 Standard and Application Benchmarks

To evaluate the performance of SGI ICE 8400 with AMD Opteron™ 6200 Series processors, the sections that follow compare detail a series of standard and application benchmarks. Results are provided for different processor models compared against SGI ICE 8400 equipped with previous-generation AMD Opteron™ 6100 series processors. Single-node results are provided for some tests while a full clustered SGI ICE 8400 system was configured to conduct SPEC MPI2007[5], OpenFOAM, and ANSYS Fluent testing. For testing, the clustered SGI ICE 8400 system was configured as follows:

- 64 compute blades (32 IP-106 compute blades and 32 IP-110 compute blades)
- 128 AMD Opteron™ 6282SE processors (16-core, 2.6 GHz) for a total of 2048 cores
- 128 GB of memory per compute blade
- SUSE Linux Enterprise Server 11SP

5 SPEComp, SPECompM, SPECompL, SPECfp, SPECint, and SPEC MPI are trademarks or registered trademarks of the Standard Performance Evaluation Corporation. Please visit www.spec.org for the latest results.

4.1 Standard Single-Node Benchmark Comparison

Single-node application performance is an important indicator for overall cluster performance. In this benchmark a system with AMD Opteron™ 6282SE processors and DDR3-1600 memory was compared against a system utilizing AMD Opteron™ 6180SE processors and DDR3-1333 memory in a variety of single-node benchmarks. As shown in Figure 5, the ICE 8400 system equipped with AMD Opteron™ 6282SE processors produced a 39% increase in Linpack performance, and a 31% increase for STREAM triad. These results speak to the intelligent overall design of the system. In particular, these performance improvements demonstrate that the increased memory bandwidth available within the system accommodates the larger number of cores per socket provided by AMD Opteron™ 6200 Series processors.

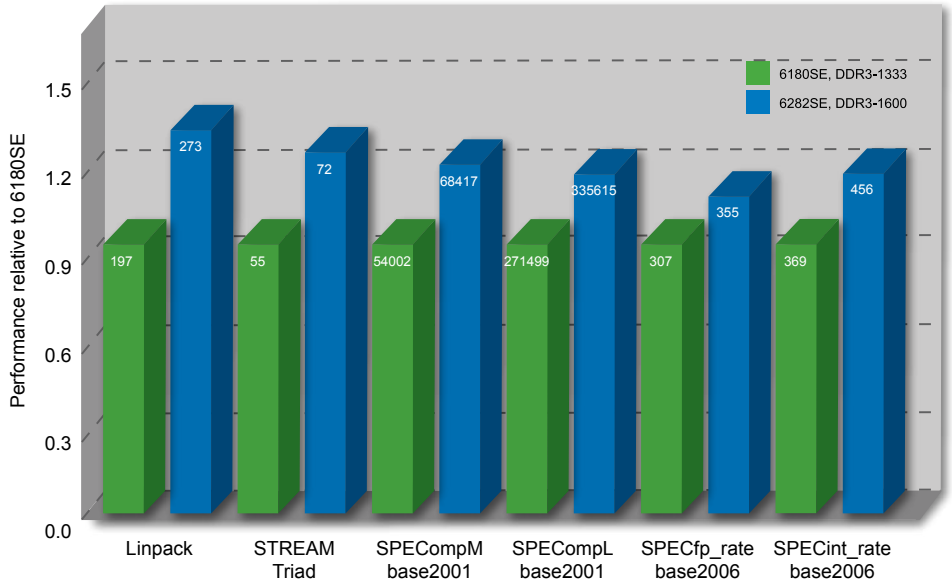

Figure 5. The single-node benchmark comparison demonstrates both scalability and balanced system design.

4.2 SPECfp_rate_base2006 and SPECint_rate_base2006

To assess the performance of ICE 8400 systems based on AMD Opteron™ 6200 Series processors, both SPECint_rate_base2006 and SPECfp_rate_base2006 tests were compared against similar systems equipped with previous-generation AMD Opteron™ 6100 Series processors. As shown in Figure 6, the tests revealed a 14% increase in floating point performance within the same physical footprint, key for data centers with restricted real estate.

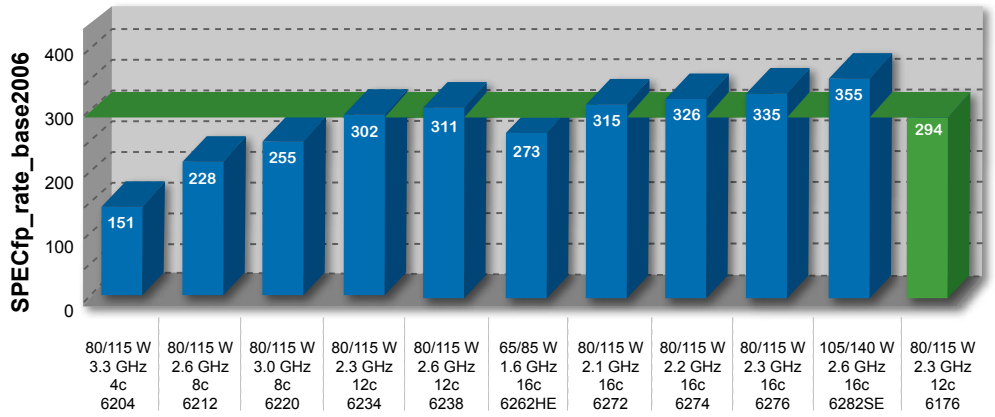

Figure 6. AMD Opteron™ 6200 SPECfp_rate_base2006 results compared with AMD Opteron™ 6176. *(Note: some SPEC CPU2006 results reflect estimates based on unpublished internal testing.)*

Shown in Figure 7, SPECint_rate_base2006 testing yielded a 19% improvement in integer throughput performance. It is important to note that this performance Improvement occurs even at the same power and thermal envelope as the tested AMD Opteron™ 6176 processor, helping organizations get more out of their existing data center environment.

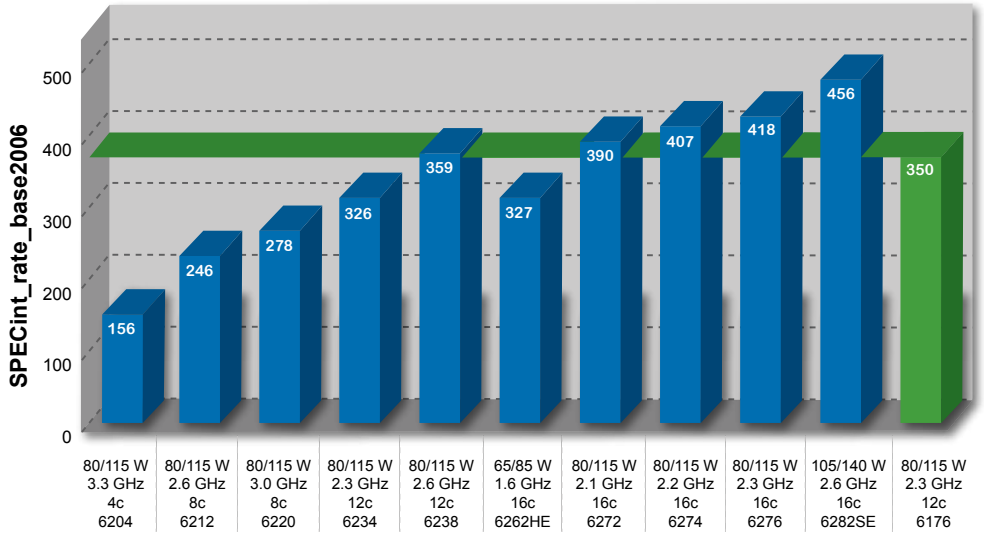

Figure 7. Multiple AMD Opteron™ 6200 Series processor SPECint_rate_base2006 results compared with AMD Opteron™ 6176 processors. *(Note: Some SPEC CPU2006 results are estimates based on unpublished internal testing.)*

4.3 SPEC MPI2007

SPEC MPI2007 is SPEC's benchmark for evaluating MPI-parallel, floating point, compute intensive performance across a wide range of cluster and MPI hardware. SGI and AMD tested a range of ICE 8400 clusters equipped with AMD Opteron™ 6200 Series processors. As shown in Figure 8, the 32-socket ICE 8400 system with AMD Opteron™ 6200 Series processors produced a 28% increase in performance when compared to a similar system equipped with AMD Opteron™ 6100 Series processors. The scalability implicit in these results clearly demonstrates that more cores equates to better real performance on SPEC MPI2007 benchmarks.

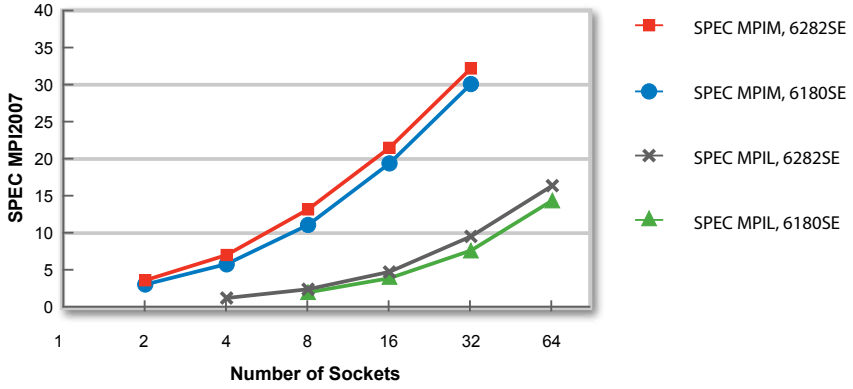

Figure 8. SGI ICE 8400 clusters equipped with AMD Opteron™ 6282SE processors provides both scalability and superior performance when compared to a similar system equipped with AMD Opteron™ 6180SE processors.

4.4 OpenFOAM

OpenFOAM is a free, open source computational fluid dynamics (CFD) packaged developed by the OpenFOAM Team at SGI, and distributed by the OpenFOAM Foundation. With a large user base across most areas of engineering and science, OpenFOAM has an extensive range of features to solve anything from complex fluid flows involving chemical reactions, turbulence, and heat transfer, to solid dynamics and electromagnetics.

In SGI testing, 1, 2, 4, 8, 16, and 32-node SGI ICE 8400 configurations were tested to compare different generations of AMD Opteron™ processors (Figure 9). The ICE 8400 system equipped with 16-core AMD Opteron™ 6282SE processors provided significant scalability running the Motorbike 11m data set over a similar configuration using 12-core AMD Opteron™ 6100SE processors.

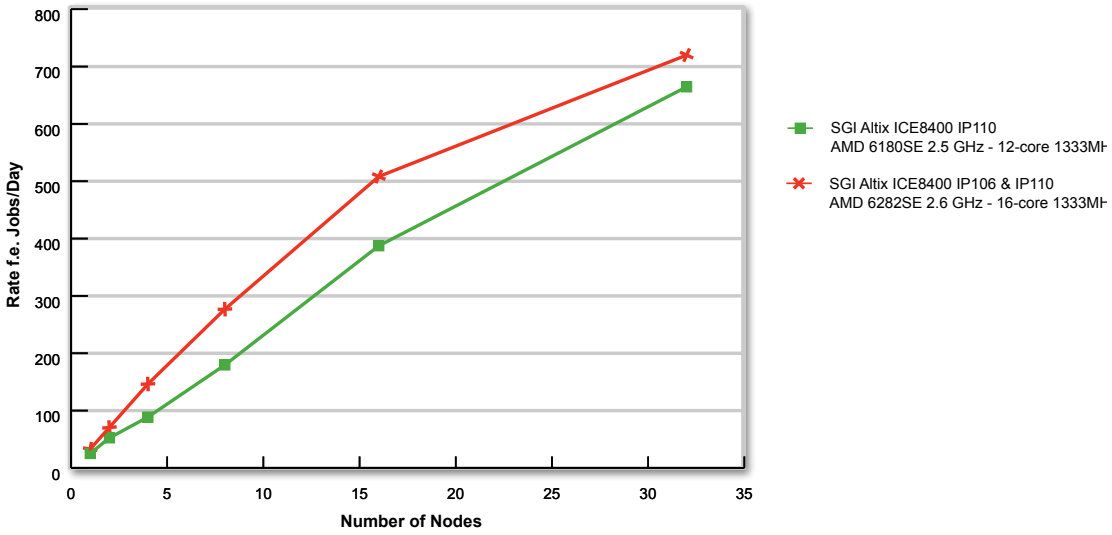

Figure 9. SGI ICE 8400 equipped with AMD Opteron™ 6282SE processors provides both scalability and superior performance when compared to a similar system equipped with AMD Opteron™ 6180SE processors running OpenFOAM on the Motorbike 11m data set.

4.5 ANSYS Fluent

ANSYS FLUENT software contains the broad physical modeling capabilities needed to model flow, turbulence, heat transfer, and reactions for industrial applications. Today, thousands of companies throughout the world benefit from use of ANSYS Fluent in applications ranging from modeling air flow over an aircraft wing to combustion in a furnace, from bubble columns to oil platforms, from blood flow to semiconductor manufacturing, and from clean room design to wastewater treatment plants. Special models have broadened the use of ANSYS Fluent, allowing it to model in-cylinder combustion, aeroacoustics, turbomachinery, and multiphase systems.

As shown in Figure 10, an SGI ICE 8400 system equipped with 16-core AMD Opteron™ 6282SE processors provides considerable scalability and performance advantage over an ICE 8400 system configured with AMD Opteron™ 6180SE processors. This result shows continued performance advantages with the addition of nodes to the problem.

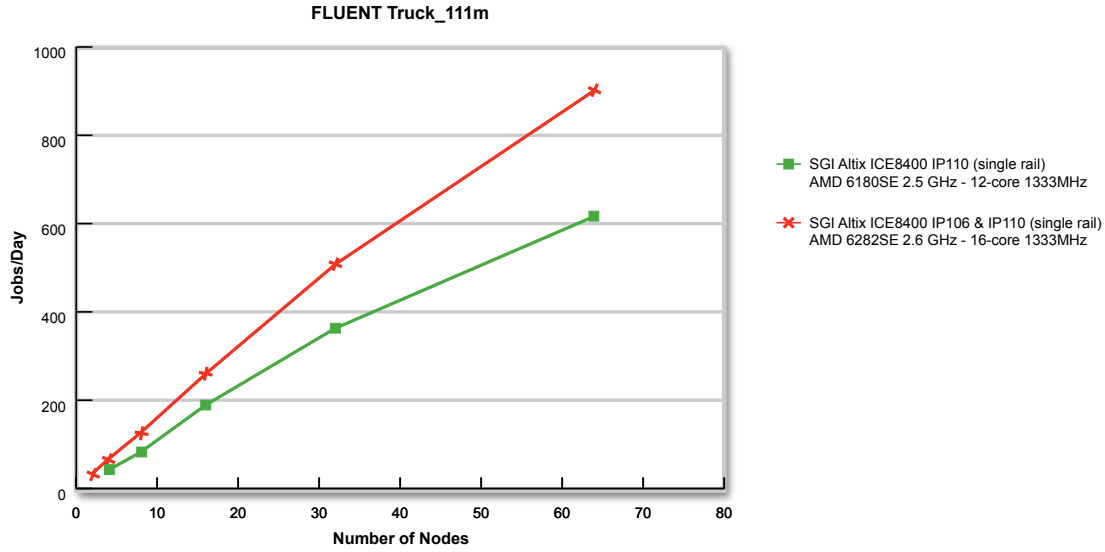

Figure 10. SGI ICE 8400 equipped with AMD Opteron™ 6282SE processors provides an increasing performance differential when compared to a similar system equipped with AMD Opteron™ 6180SE processors running ANSYS Fluent on the Truck_111m data set.

5.0 Conclusion

SGI ICE 8400 combined with AMD Opteron™ 6200 Series processors represents a unique combination of innovative technologies that can pay real dividends in terms of density, performance and scalability, all in a manageable system that can be deployed rapidly. SGI ICE 8400 represents a proven and fully integrated system that offers extreme processing efficiency and flexible choices in terms of storage and interconnect topology. With up to 16 cores, AMD Opteron™ 6200 Series processors add considerable computational prowess, along with the density and environmental advantages to make these systems easy to deploy in today's constrained data center environments.

6.0 References

For more information on SGI ICE 8400 systems, please visit http://www.sgi.com/products/servers/ice/8400

For more information on AMD Opteron™ 6200 Series processors, please visit: http://www.amd.com/opteron

For AMD performance benchmarks, please visit: http://www.amd.com/benchmarks

To sign up to test drive your applications, please visit: http://www.sgi.com/amd

Global Sales and Support: sgi.com/global

© 2012 SGI. SGI, Altix, NUMAlink, XIO, SGI Fullcare, SGI Fullcare Express, and SGI FullExpress 7x24 are registered trademarks or trademarks of Silicon Graphics International Corp. or its subsidiaries in the United States and/or other countries. All other trademarks are property of their respective holders. 17072012 4385