INNOVATION FOR RESULTS™

White Paper

SGI® Multi-Discipline Simulation Solutions

Table of Contents

1.0 Building a Comprehensive HPC Solution for the CAE Environment

A growing number of manufacturers facing heightened consumer expectations and increasingly stringent regulatory requirements now turn to advanced simulation technologies– including a broad range of computer-aided engineering (CAE) applications - to gain a competitive advantage. The use of this technology in product development can deliver a critical competitive edge by reducing design cycle time while lowering costs and improving product quality and reliability.

Manufacturing involves intensive collaboration between many disciplines. CAE applications that address those disciplines - such as computational fluid dynamics (CFD), finite element analysis (FEA) and kinematics used to evaluate performance, aerodynamics, structural fatigue, crash worthiness and noise, vibration and harshness characteristics - make it possible to predict system behavior accurately. To realize the full benefits of these multi-disciplinary simulations, however, the manufacturer will need a high performance computing (HPC) infrastructure optimized to leverage their power. This will include a balanced HPC architecture that meets the needs of both capability and capacity, ready structured access to all simulation data from workgroup level to enterprise level over the complete product lifecycle and fully integrated resources for post-processing of simulation results.

2.0 The Multi-Disciplinary World of CAE

Because HPC applications – and CAE applications in particular – have unique processing requirements and workloads, a multi-workflow architecture for high performance computing is in order.

A number of challenges exist in developing an infrastructure to maximize CAE for product design scientific simulations. The multi-disciplinary environment in which the applications operate involves substantial interdependency and interaction between the various disciplines. Additionally, different physics-based solution procedures and disparate models and tools used within individual disciplines must be applied for adequate prediction of design behavior. Also, iterative design and analysis process results in several hundreds of multi-disciplinary simulations for a converged solution.

Figure 1: Hybrid Architecture

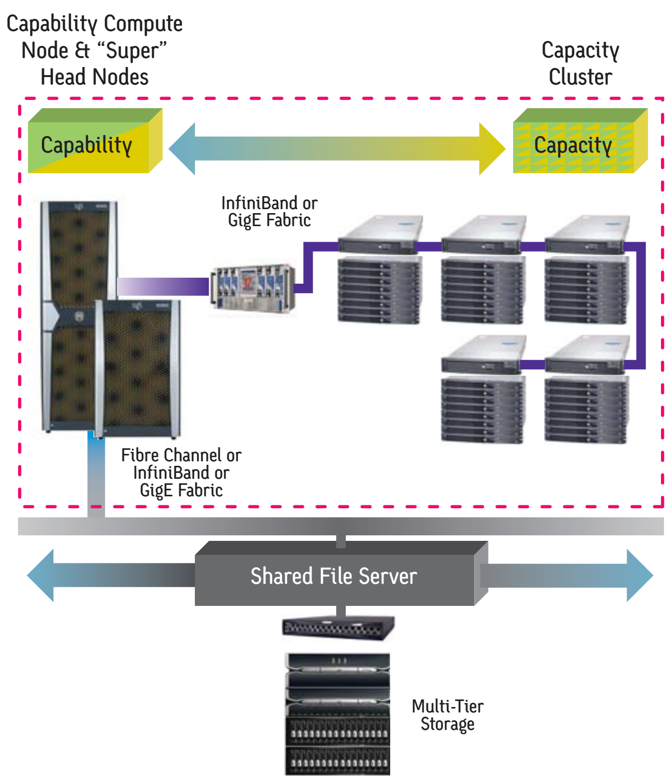

Unfortunately, the HPC environment that many industries face as they work to solve these problems is made up of geographically-distributed specialty teams working on various aspects of a complete system, including the subsystem and disciplinary levels. Additionally, manufacturers are challenged by a heterogeneous computing environment based on multiple system architectures and operating systems to meet the individual discipline application needs. That results in increased data transfer between host systems. The situation is further complicated by the explosion of data resulting from the use of high fidelity models and the increase in numbers of simulations. If, as is often the case, integration between islands of systems and disciplines and structured access to data are lacking at departmental levels and enterprise-wide, the circumstances may be worsened.

This less-than-ideal HPC environment frequently results in post-processing of very large disciplinary analysis models and performance and resource limitations. Workflow bottlenecks related to data flow from conceptual design to detail computer-aided design (CAD) modeling and then to CAE analysis and model verification in closing the design analysis loop associated with iterative design improvement processes can also be a problem.

One of the primary challenges identified in optimizing CAE applications in multi-disciplinary simulations is designing a balanced HPC architecture to address two distinct and competing workloads: capability and capacity computing.

Capability workloads solve large complex problems and require the most powerful processing to minimize the time-to-solution duration. Typically, this would be dedicated to running no more than a few problems simultaneously. These workloads benefit from big cache, big memory, big disk storage and high-speed I/O.

Capacity workloads, on the other hand, involve a large number of parallel jobs in which each job has moderate computational requirements. Less data sharing between applications and more parallel application threads – smaller, parallel jobs – set these workloads apart from capability workloads.

To accommodate the multiple, diverse disciplines needed to accurately evaluate product behavior and meet product quality requirements, manufacturers need an HPC solution that addresses both capability and capacity workloads. The balanced HPC architecture will also make it possible to increase model fidelity and produce accurate boundary conditions, statistical model verification and multi-disciplinary analysis and design optimization.

Another requirement for meeting infrastructure challenges lies in addressing multi-disciplinary simulations that involve problems with varying physics that can't be solved in a single package. This is necessary to overcome the lack of integration among simulation technologies and software that have evolved as "islands" of individual discipline technologies.

Gaining structured access to product and simulation data enterprise-wide, too, has been a problem in multi-disciplinary simulations. Manufacturers need to be able to link data and processes that involve mixed architectures and heterogeneous data environments across the extended enterprise. Equally as important is the manufacturer's capability for data management over the entire product lifecycle.

3.0 Optimizing the HPC Architecture for CAE

How then, can a manufacturer design an HPC infrastructure for multi-disciplinary projects?

First, organizations need to recognize that their HPC infrastructure will need to accommodate a number of product disciplines. Structures such as ANSYS, Nastran and ABAQUS, impact such as LS-DYNA, and CFD such as CFX, FLUENT and STAR-CD are all considerations.

Secondly, they must include a mix of different computer architectures, ranging from large SMPs to small node clusters. This will make it possible to support the hybrid workloads of capability and capacity-centered computing and accommodate both parallel scalability and I/O throughput intensive jobs for improved time-to-solution.

Additionally, the infrastructure must be designed to handle a dramatic increase in the number of simulations as a result of increasing boundary conditions, statistical model verification and multi-disciplinary optimization. This will require centralized, unified storage to minimize the data transfer between different host systems. It will also help meet the need for consistent data management providing for fast, structured access to all simulation data across the enterprise.

Finally, a successful multi-disciplinary design project will need specialty teams working on various disciplines. Though likely to be geographically distributed, the teams must be able to work together effectively to process integration frameworks for seamless integration of islands of systems and disciplines with coupling tools and networked visualization.

4.0 One Size Does Not Fit All

The numerical characteristics of applications in various CAE disciplines will determine what type of HPC architecture would best suit the needs of the codes involved. Recognizing this, SGI has incorporated leading edge technology in a unified platform to provide a flexible solution that targets specific workflow requirements. In this type of architecture, unified system software and a common storage infrastructure enable a single, multi-platform workflow.

SGI solutions address the needs of all CAE customers, from those needing the most powerful system to those needing a cost effective platform to run a large number of simultaneous jobs with moderate computational requirements.

Implicit Finite Element Analysis (IFEA), involving the evaluation of mechanical components and systems for structural strength, vibration and heat transfer, is one such example. In applications requiring large amounts of memory and non-linear problems, as well as significant capabilities for handling noise, vibration and harshness (NVH) the SGI® Altix® 4700/450 meets customers' CAE needs. If, on the other hand, the application calls for a system offering static and non-linear analysis, moderate eigensolutions for average problems where I/O needs are not demanding, the SGI® Altix® XE would be best suited.

Explicit FEA (EFEA), which performs evaluation of structural failure owing to high frequency load conditions such as impact, or high speed collision with rapid change over time and involving highly non-linear effects, has its own unique computing requirements. EFEA performance is optimal on clusters with small node sizes and has applications. For that reason, the Altix XE would provide the ideal solution.

In the field of CFD, for commercial and government off-the-shelf (COTS and GOTS) applications involving the evaluation of flow fields such as liquids and gases, the system will need to appropriately address transient and intransient steady state flow problems. For a class of large scale transient CFD problems characterized by significant amounts of data transfer, the Altix 4700/450 delivers the most compelling performance. On the other hand, for a large chunk of the daily CFD workload, especially involving steady state fluid flow, the Altix XE cluster system is the ideal workhorse,as the computational workload is easily balanced across multiple processors.

In Computational Electromagnetics (CEM) - which involves cross section signature analysis of planes, ships and ground vehicles, as well as scattering and target identification and RF performance and isolation - frequency is key. For CEM involving low frequency, the Altix 4700/450 is suitable. For high frequency, the Altix XE is ideal.

Figure 2: MDS Rack

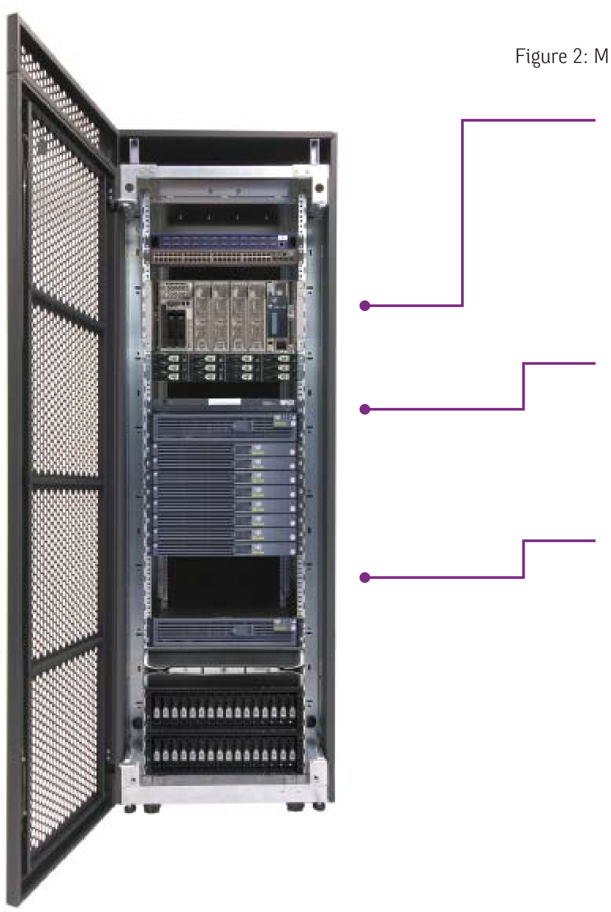

Altix 450 and 4700 for High-Capability Shared-Memory Workflows

SGI Altix 450/4700 servers and supercomputers deliver industry-leading performance, scalability, and versatility with SGI NUMAflex™ shared-memory architecture and a revolutionary blade-based design for perfect system right-sizing. For applications requiring large amounts of memory and high-speed I/O, such as is the case with NVH, non-linear, and large scale transient fluid flow problems, the SGI Altix 450 and 4700 satisfy the most demanding CAE needs.

Altix XE for High-Capacity Clustered and Distributed Memory Workflows

SGI Altix XE servers and clusters offer superior price-performance, compute density and energy efficiency for capacity-oriented worklows. Scalable Altix XE head and compute nodes using InfiniBand interconnects are particularly well suited for applications such as crash analysis and CFD.

Flexible and Comprehensive Data Management

SGI InfiniteStorage solutions offer a full line of state-of-the-art disk storage systems designed for data-intensive CAE environments. SGI MDS solutions deploy performance-oriented primary and capacity-oriented secondary storage subsystems, well suited for iterative design environments that rely on simulation to improve quality, reduce costs and shorten time to market. SGI® Data Migration Facility (DMF) automates data migration for the most cost-effective and highest possible capacity utilization across all storage.

Customers seeking CAE optimization are not limited to one SGI system or another, however. In fact, the SGI Altix 4700/450, with its high capability shared memory workflows and the SGI Altix XE with its high capacity clustered and distributed memory workflows, can be configured as a factory integrated solution. The customizable solution stacks can then be targeted to specific solutions.

5.0 Customizable and Scalable Hybrid Solution Stacks

While the Altix 4700/450 and Altix XE are key components in the CAE solution, it is important to note that SGI integrates a customized combination of computing and storage platforms with system management tools for a total, comprehensive solution. This hybrid solution stack includes SGI InfiniteStorage solutions and network interfaces. For CAE end users, the hybrid solution stack provides the best of both worlds with an architecture that ensures the best possible application level performance without over-subscribing the hardware. Its building blocks are configured based on application specifications and appropriate storage components. Additionally, the hybrid solution stacks address IT management, providing maximum efficiency in the IT infrastructure and a tightly integrated environment through CXFS® (Clustered Extended File System) and DMF capabilities. The hybrid solution stack makes it possible for CAE users to meet price performance targets without compromising on absolute performance needs by selecting just the right amount of appropriate compute blades to meet the needs of the workload.

6.0 The Total Package

SGI Altix and Altix XE clusters can be seamlessly integrated using SGI® InfiniteStorage, enabling an efficient data management scheme for both capability and capacity workflows, allowing customers to reduce costs and speed overall product development. Scalable storage is critical in processing the dynamic, large data sets that make it possible for teams of CAE engineers to access and manage simulation data. SGI's storage solution can scale in terms of bandwidth, capacity and connectivity. Its primary and secondary storage subsystems are well suited to iterative design environments that rely on simulation or virtual prototyping to reduce costs and shorten time to market. High speed primary storage shortens critical process, and cost effective secondary storage handles lower priority, less frequently accessed data.

Additionally, SGI storage solutions provide comprehensive data management through a full line of disk storage systems designed for data intensive CAE design environments. The data-centric architecture delivers optimal efficiency for storing, organizing, accessing and managing data, and provides high-speed, shared access that eliminates the bottlenecks that can hinder virtual prototyping and visualization operations so critical to CAE.

Figure 3: Complete Solutions for the Enterprise with Centralized Data Management

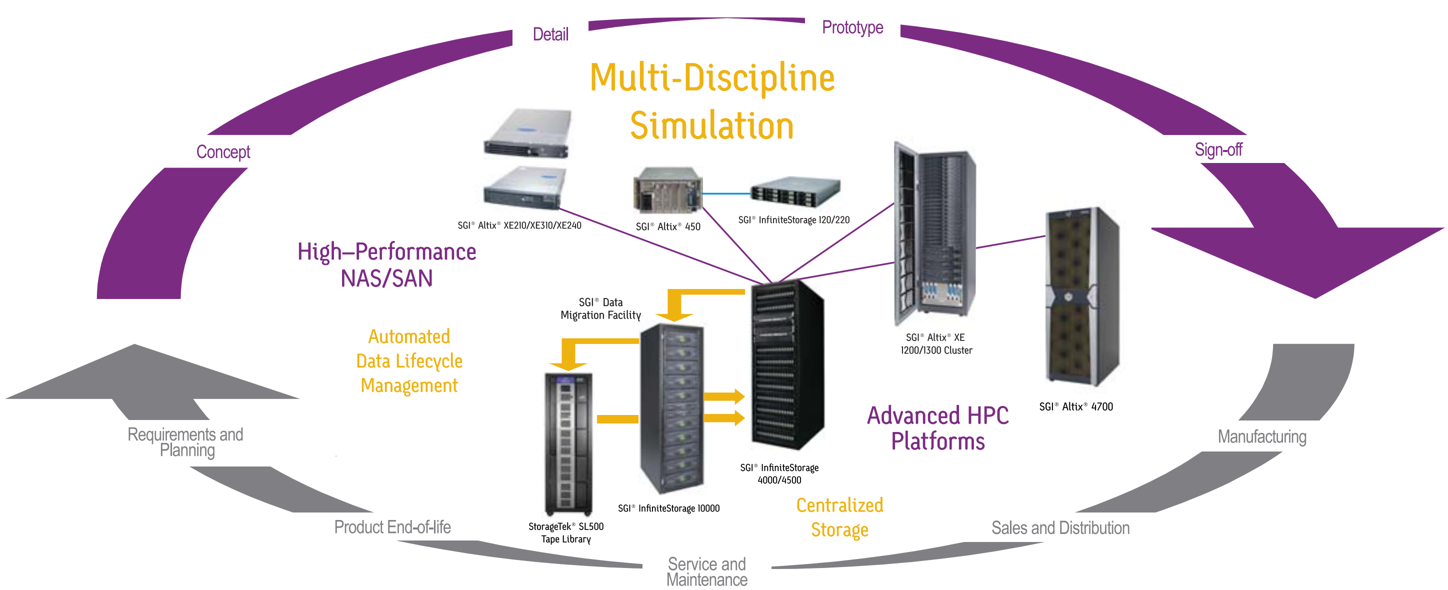

The SGI solution is not made up of technology alone. A key element in providing complete solutions for CAE workflow is the company's application and domain expertise. SGI's ability to customize the workflow with the right mix of Altix XE and Altix servers and InfiniteStorage is instrumental in enabling customers to maximize their IT investments. Its global Professional Services team stands ready to provide a complete suite of services to cover every aspect of CAE solution deployment. Additionally, SGI has developed a strong ISV ecosystem. Customers benefit from guaranteed multi-year ongoing support and certification of applications on SGI systems as well as specific initiatives aimed at code optimizations and integration of SGI value-add in the middleware layer. As a result of its strong ISV relationships, SGI is able to escalate problem resolution and to establish joint initiatives with ISVs to address specific customer challenges.

7.0 Conclusion

CAE workflow environments usually involve multiple applications from multiple disciplines, each with differing numerical characteristics. No single system architecture can handle all these applications and be best in all. SGI now offers two complementary architectures leveraged for appropriate applications, SGI Altix servers and SGI® Altix® Clusters, to address the entire CAE workflow. This offers the customer the flexibility to choose a system architecture based on the numerical attributes of the codes.

SGI meets the critical need for data management and storage with a storage solution that unifies data storage, shared access, movement and management across all systems. With its established application expertise and domain knowledge, SGI has emerged as a single provider of all compute and storage solutions for the entire CAE workflow.

Corporate Office
1140 E. Arques Avenue
Sunnyvale, CA 94085
(650) 960- 1980
www.sgi.com

North America +1 800.800.7441
Latin America +55 11.5185.2860
Europe +44 118.912.7500
Japan +81 3.5488.1811
Asia Pacific +1 650.933.3000

3981 [03.2007] **J15207**