White Paper

Extending the High Performance Computing Workflow with SGI® Altix® XE Clusters Running Microsoft® Windows Compute Cluster Server 2003

Table of Contents

1.0 Overview

Clusters offer an economical approach to high performance computing (HPC). As such, they have been the main force driving the recent rapid expansion of HPC, with an installation growth rate of approximately 50% per year over the last five years. The ongoing acceleration of higher performance for lower cost is reducing HPC cost-of-entry and helping to bring the benefits of parallel computing to an ever-greater range of applications and organizations. Clusters' inherent modularity means that businesses can start small and build their HPC environments as their need grows.

It is not enough, though, that the price of raw computing power continues to diminish rapidly. Cluster computing often comes at a cost. To take advantage of inexpensive HPC clusters, companies must also somehow acquire the knowledge required to design, manage, and run their applications in a clustered environment. SGI® Altix® XE clusters and Microsoft® Windows Compute Cluster Server 2003 (WCCS) software work together to lower the knowledge bar to match the cost-of-entry bar. They extend HPC to a whole new range of applications, while bringing new levels of ease-of-use and value to existing HPC application areas.

SGI Altix XE clusters provide superior price performance and energy efficiency for mainstream high performance computing, with flexible solutions designed to handle a wide range of application areas, from scientific research and computer-aided engineering (CAE) to finance and digital content creation (DCC). SGI Altix XE combines the powerful Intel® Dual-Core and Quad-Core Xeon® processor-based architecture with SGI's expertise in designing and delivering some of the most advanced HPC systems available today. The result is breakthrough value and ease-of-use with no-compromise performance.

WCCS on SGI Altix XE adds the familiarity and usability of Windows to the best of value-focused HPC, at the same time enabling close integration with a wide range of Microsoft® products and associated development and desktop analysis tools. WCCS includes a complete, integrated software stack for developing, managing, and running HPC applications within a Windows environment.

The winning combination of SGI Altix XE clusters and WCCS means users can maintain a results-based focus with their HPC projects while taking advantage of their prior experience with familiar tools, instead of having to devote valuable research time to mastering IT-level intricacies of managing cluster environments. SGI Altix XE and WCCS together bring an unbeatable combination of value, performance, and ease-of-use to the world of mainstream high performance computing.

2.0 SGI Altix XE Clusters: Cutting-Edge Solutions at an Affordable Price Point

SGI's Altix XE line of servers and clusters offers exceptional value for high performance computing. SGI Altix XE incorporates the latest in Intel's advanced Dual-Core and Quad-Core Xeon processor architecture into dense and energy-efficient packages.

A cluster consists of a set of tightly integrated servers, with at least one head node and multiple compute nodes per cluster. The head node handles overall cluster administration functions. Actual application processing occurs in parallel across the set of compute nodes. Depending on the need of the application, the number of compute nodes in a cluster can range from just a few to hundreds or even thousands. SGI Altix XE clusters are easy to scale, enabling users to design a system that meets their specific, current needs, while reserving the potential to scale the system in the future as processing needs grow or change.

Utilizing SGI Altix XE servers for their compute and head nodes, SGI Altix XE clusters provide flexibility in design to meet the needs of a wide range of HPC applications, with value-focused emphasis on performance, density, efficiency, and connectivity. SGI Altix XE servers include advanced features such as a super-fast 1600MHz front-side bus, 64GB of memory per compute node, and an ultra-dense architecture that packs up to two eight-core nodes in a slim 1U form factor (Figure 1).

Figure 1. SGI Altix XE310: Two Nodes on a Single Board

They use fully buffered DDR2 memory and support 20GB/second InfiniBand and/or Gigabit Ethernet interconnects. The SGI Altix XE servers are available for purchase as high performing standalone workgroup servers, but their design is particularly well suited for use as nodes in cluster configurations. SGI Altix XE clusters incorporate the servers in a variety of configurations and performance options.

SGI Altix XE servers integrate easily with SGI InfiniteStorage, a comprehensive storage infrastructure that is used to store, manage, and access data across heterogeneous data sources, with up to 224TB of stored data on a single storage appliance.

SGI Altix XE clusters with WCCS are delivered with a complete, integrated cluster solution stack from Microsoft, including Compute Cluster Administrator, Compute Cluster Job Scheduler, and MS-MPI. Additionally, WCCS takes advantage of existing Windows infrastructure tools such as Active Directory®.

SGI Altix XE clusters also offer the option of combining WCCS and Linux on the same cluster, providing maximum flexibility for shops running HPC applications in both Linux and Windows environments. Depending on your application needs, you can either dedicate the entire cluster alternately to Linux or WCCS, or run Linux and WCCS applications simultaneously on separate compute nodes within the cluster.

Two WCCS/Linux cluster combinations are available on SGI Altix XE: dual-boot and "split-brain". Both options require separate WCCS and Linux head nodes. The difference lies in how the compute nodes are configured. In the dual-boot configuration, each compute node includes both Linux and WCCS, residing on separate partitions. Depending on the cluster's application load, you can boot some or all compute nodes as WCCS or Linux. When the application load changes, changing the mix of WCCS and Linux compute nodes is a simple matter of rebooting the nodes.

With the split-brain approach, a single operating system resides on each compute node. You decide when ordering the cluster what combination works best for your application needs. For example, if the majority of your applications run on WCCS, you might choose to configure a 24-compute-node system with 16 WCCS nodes and 8 Linux nodes. The split-brain option is simpler to set up and configure, but it doesn't offer quite the same degree of flexibility as the dual-boot option. In either case, SGI will deliver SGI Altix XE clusters to your data center pre-loaded with whatever options you require.

For mixed workflow environments, SGI Altix XE clusters can be deployed in combination with SGI® Altix® servers and supercomputers or SGI® Altix® ICE advanced blade servers, to optimally address diverse processing requirements. This hybrid compute environment (Figure 2) offers the best of capacity and capability compute solutions, with an ability to scale up and/or scale out as needed to meet current and future processing needs. SGI Altix XE clusters lead the value side of the SGI server product line, offering high performance solutions with easy scalability at an attractive price point.

2.1 The SGI Altix XE Line of Clusters

SGI Altix XE clusters are custom configured and pre-built, providing maximum flexibility to handle a wide diversity of performance computing needs.

The SGI Altix XE1200 cluster configuration uses SGI® Altix® XE250 compute nodes and offers advanced extensibility, with a rich set of expansion and I/O options to address the broadest range of compute requirements.

The SGI® Altix® XE1300 cluster configuration, with SGI® Altix® XE310 or SGI® Altix® XE320 compute nodes built on the innovative Atoka board design, supports up to 16 processor cores in a 1U blade, delivering industry-leading performance density with cost efficiency.

Both SGI Altix XE1200 and SGI Altix XE1300 clusters incorporate SGI Altix XE250 cluster head nodes, for maximum extensibility and power.

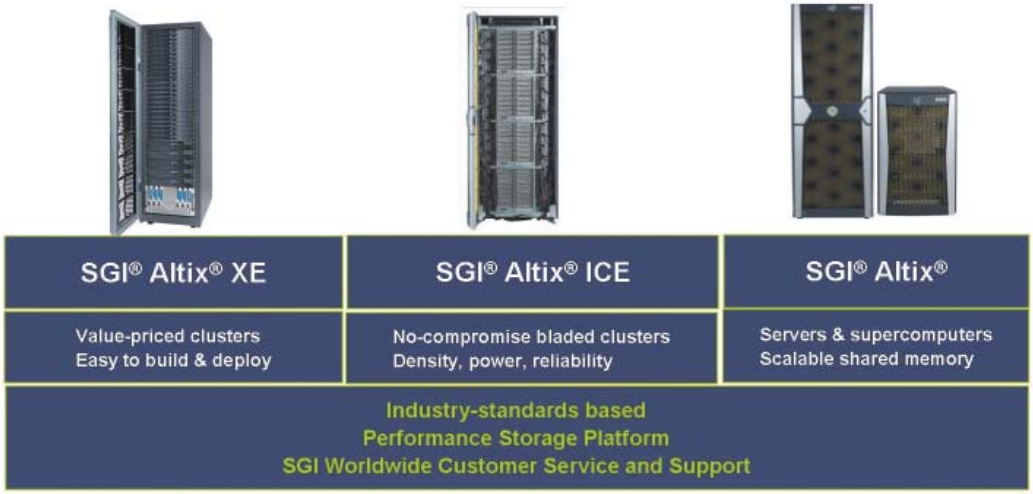

Figure 2. Addressing the High Performance Workflow Continuum

	Altix XE310	Altix XE320	Altix XE250
Chassis	1U	1U	2U
Processors	Up to four Quad-Core Intel® Xeon® Processors, 5400 Series (2 per node) - Front Side Bus: 1333 MHz - L2 Cache: 12MB	Up to four Dual or Quad-Core Intel® Xeon® Processors, 5200 or 5400 Series (2 per node) • Front Side Bus: 1600 or 1333 MHz • L2 Cache: 6MB for 5200 series, 12MB for 5400 series	Up to two Dual or Quad-Core Intel® Xeon® Processors, 5200 or 5400 Series • Front Side Bus: 1600 or 1333 MHz • L2 Cache: 6MB for 5200 series, 12MB for 5400 series
Internal Storage	Four SATA drive bays (2 per node) Four SAS drives available via optional PCIe card Optional RAID 0, 1		Eight SATA/SAS drive bays 1 x DVD-RW drive RAID 0, 1, 5 (optional), or 10
Memory	64 GB DDR2 800 or 667 MHz FBDIMM memory (32GB per node) Supports memory sparing and mirroring		64 GB DDR2 800 or 667 MHz FBDIMM memory Supports memory sparing and mirroring
Add-in Card Support	2 x PCIe x8 (1 per node)	2 x PCIe x16 (1 per node)	2 x PCIe x8 gen 2 1 x PCIe x8 gen 1 1 x PCIe x4 gen 1 1 x PCI-x 133/100
Benefits	Dense compute node Low-cost	Dense compute node, with enhanced processing capabilities	Powerful cluster head node I/O rich, extensibility, redundancy

Table 1. SGI Altix XE Server Characteristics

3.0 WCCS: Integrating SGI Altix XE Clusters into the Windows Environment

WCCS extends the SGI Altix XE value proposition by providing Windows-level ease-of-use and system integration capabilities. WCCS is a complete HPC software stack, built on a core of Windows Server 2003 x64 edition, with numerous enhancements to provide HPC development, deployment and management capabilities. It brings to high performance computing all the benefits of a comprehensive Windows infrastructure, such as integration with Active Directory and access to a wide range of third party applications and development tools, including Microsoft Visual Studio®. WCCS couples seamless system integration with tools and capabilities that provide users full control over their HPC environment within a familiar Windows interface.

3.1 A Comprehensive HPC Software Stack

WCCS integrates all the tools and services needed to manage HPC clusters and jobs, providing a quick and seamless route to HPC success. It consists of two main components:

- Windows Server 2003, Compute Cluster Edition – an HPC edition of x64-based Windows Server 2003 priced and licensed for HPC.
- Compute Cluster Pack – a set of tools and utilities for managing HPC clusters, including:
  - Compute Cluster Administrator
  - Compute Cluster Job Scheduler, with Compute Cluster Job Manager
  - MS-MPI

The Compute Cluster Pack features a complete set of tools and services for configuring the cluster, managing the nodes, and creating and submitting jobs. Besides a comprehensive Windows-based GUI, made even simpler through frequent access to wizards for performing configuration tasks, most functions are also accessible through a command line interface (CLI) or an Application Programming Interface (API).

The Compute Cluster Administrator is a GUI for configuring the cluster, managing nodes, and monitoring cluster activity and health. The Compute Cluster Job Manager handles job creation, submission, and monitoring.

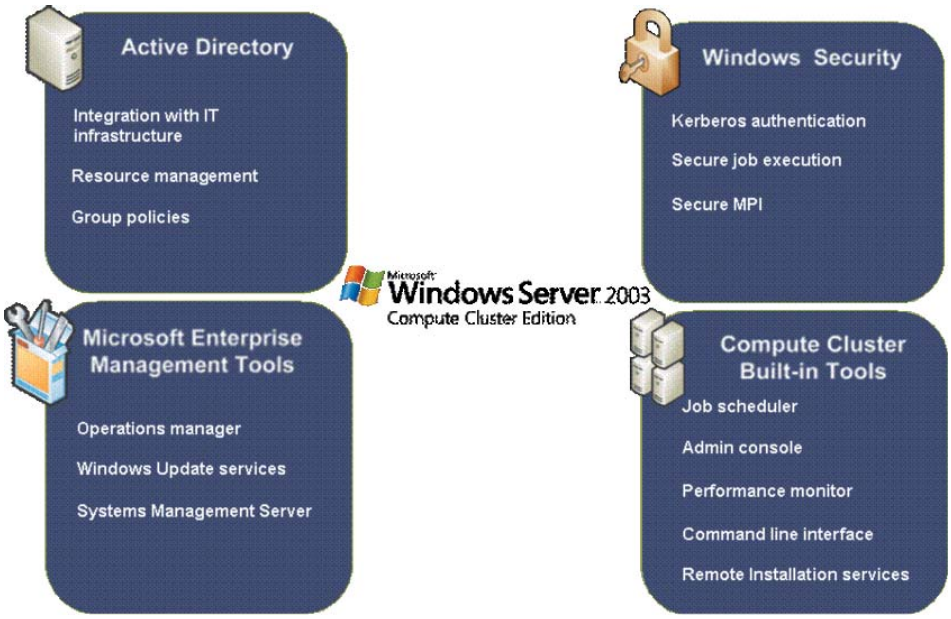

Image reprinted with permission from Microsoft Corporation.

Figure 3. CCS Comprehensive Solution.

Message Passing Interface (MPI) is a standard specification for message passing, designed specifically for HPC parallel applications. MS-MPI is Microsoft's version of the widely used Argonne National Labs Open Source MPICH2 Reference Implementation. Microsoft Visual Studio 2005 includes a parallel debugger that works with MS-MPI, enabling developers to launch their HPC applications on multiple compute nodes and then pause and examine program variables on each node.

3.2 Integration with Existing Windows Infrastructure

SGI Altix XE clusters running WCCS integrate easily into the existing Windows-based IT infrastructure. Of key importance, WCCS uses the Active Directory service to manage directory, authentication, and security services. With Active Directory, a network user requires only a single logon process to access all permitted services across the network, including those on SGI Altix XE cluster head and compute nodes.

WCCS also uses the Microsoft Management Console (MMC) for extensible snap-ins and integration with the Microsoft Operations Manager 2005 (MOM), Remote Installation Services (RIS), and the Microsoft Systems Management Server (SMS).

3.3 An HPC Development Environment with Leverage

WCCS developers can create HPC applications with Visual Studio 2005, which incorporates a parallel debugger and support for MPI. This not only means that developers can leverage previous experience with Visual Studio, but it also opens up the possibility of porting existing, non-parallel Visual Studio applications to SGI Altix XE clusters. WCCS can also interact with other Microsoft products, such as Microsoft Office Excel® – of key importance to the burgeoning use of HPC in financial applications. Finally, SGI Altix XE users can take advantage of the large number of third-party software tools available on WCCS in a number of application areas, including manufacturing, life sciences, geological sciences, and financial services.

4.0 SGI Altix XE Clusters Running WCCS: Synergy that Works … For You

With SGI Altix XE, SGI maintains its leadership in offering performance with value. SGI Altix XE clusters meet the needs of a wide range of high performance computing requirements by offering the latest in design innovations – at an unbeatable value point, in both initial and long-term costs. WCCS supplements the SGI Altix XE value proposition by providing Windows-level usability and integration throughout the development and implementation lifecycle.

4.1 Providing a Path to Quick Deployment

SGI is renowned for delivering products with "power-up-and-go" factory pre-building and testing, ensuring immediate customer productivity and easing the IT burden. SGI Altix XE clusters arrive at the customer site pre-built, tested, and ready-to-go. WCCS further minimizes time-to-productivity by providing an automated, task-based interface for setting up and deploying the clusters. The comprehensive set-up process assists with networking, remote installation services, node management, and cluster security.

4.2 Simplifying Manageability

SGI factory pre-building means that you can manage the SGI Altix XE cluster as a single unit, rather than starting from a set of building-block servers requiring significant expenditure in time and effort to integrate and prepare for use. With onboard InfiniBand and Gigabit Ethernet, the design of SGI Altix XE1300 reduces significantly the number of connecting cables and thus the complexity of managing or scaling the cluster hardware.

On the software side, the WCCS management tools, accessible through both user interface and command line, greatly simplify the management of SGI Altix XE cluster nodes. WCCS also makes it easy to schedule jobs on the cluster, with the MS Job Scheduler providing an interface akin to that used for print queue management. The familiar Windows interface ensures that administrators and developers can build on their experience with other Windows products to gain quick mastery over the WCCS toolset.

4.3 Leveraging Existing Infrastructure

The WCCS interface provides seamless integration with existing Windows infrastructures. WCCS integrates with Active Directory for easy authorization and authentication. WCCS also draws upon other Windows capabilities and tools, including Remote Installation Services (RIS), Microsoft Systems Management Server (SMS), Microsoft Operations Manager (MOM), and Microsoft Management Console (MMC).

SGI Altix XE clusters leverage your investment in data center infrastructure by reducing costs in a number of areas. The high performance, small footprint SGI Altix XE310 design, packing up to 16 core nodes in a 1U form factor, provides a level of performance density that ensures efficient use of precious data center real estate. Furthermore, SGI's industry-leading power and cooling technologies work to reduce air conditioning and other power-related costs.

4.4 An Advanced HPC Environment

The power, flexibility, and scalability of SGI Altix XE enables customers to design cluster systems that exactly meet current HPC development needs, while maintaining the ability to adjust and scale for future applications. And SGI's industry-leading Professional Services and Support organizations, with deep technology and industry experience, are always ready to provide the expertise to meet any scale-out challenge.

With SGI Altix XE running WCCS, developers can create HPC applications with the number one integrated development environment (IDE) – Microsoft Visual Studio. Visual Studio 2005 incorporates full support for HPC applications, including a parallel debugger, so developers can easily leverage existing skills as they begin to design for HPC. WCCS's integrated MPI layer also eases the effort of porting existing parallel applications. In addition, numerous third-party applications have been written for WCCS, covering a broad range of HPC needs and further minimizing customer development costs.

4.5 The Bottom Line: High Performance with Low Total Cost of Ownership

Low total cost of ownership (TCO) starts with the initial server cost, and SGI Altix XE delivers leading value in compute dense configurations with industry-leading price/performance. SGI Altix XE extends that value through innovations on a number of fronts besides raw price/performance. The dense board design, combined with a choice of onboard InfiniBand or Gigabit Ethernet, reduces the number of interconnect cards and cables, thereby reducing component costs while simultaneously increasing system reliability and manageability.

Innovative high-density compute nodes mean that SGI Altix XE clusters minimize use of valuable data center real estate. SGI Altix XE clusters also benefit from SGI's unmatched experience in power and cooling efficiencies, enabling customers to build and deploy SGI Altix XE clusters with maximum energy efficiency. Finally, SGI's industry-leading Professional Services and Support teams, combined with factory pre-building, ensure that customers experience custom-tailored solutions that are both highly reliable and simple to deploy and scale.

WCCS further promotes low TCO by helping to reduce IT and development costs. WCCS fits right into existing Windows infrastructure, eliminating any need for a large outlay in new system software capabilities. Developers can continue to work with their familiar Windows-based development tools, like Visual Studio, easing their transition to HPC. And the Windows ecosystem ensures that plenty of third-party tools are available to help shorten time-to-deployment.

5.0 Key Application Areas

SGI Altix XE running WCCS makes sense as a cost-effective, manageable solution for any application area that depends on fast processing of large quantities of complex, parallel computations. The SGI Professional Services organization offers deep and long-standing expertise in a wide range of HPC domains, including CAE, digital content creation, financial services, and the sciences. SGI's experience in these markets is bolstered by its close collaboration with key independent software vendors (ISVs). In addition, due to WCCS's integral position in the Windows ecosystem, there are certain application areas where SGI Altix XE running WCCS represents an extra-ordinarily attractive solution.

WCCS on SGI Altix XE offers particular advantage to industries such as financial services, oil and gas, and life sciences, where Excel frequently serves as the critical tool for performing highly complex data analysis and modeling calculations. Several key features of Microsoft Office Excel 2007 enable powerful integration possibilities with WCCS on SGI Altix XE. Excel 2007 provides a host of new and enhanced capabilities from earlier versions of Excel, including a massive increase in spreadsheet size limits to 16,384 columns by 1,048,576 rows, expanded addressable memory, and a multi-threaded calculation engine that enables it to accelerate calculation performance by operating across parallel threads in an HPC environment.

In conjunction with WCCS on SGI Altix XE and Excel Services – a component of Microsoft Office SharePoint® 2007 that provides a server-based version of Excel – Excel users can now take full advantage of the performance gains of HPC. Through offloading of processing from the desktop to an SGI Altix XE cluster, complex spreadsheet calculations can run in parallel, and users get results faster. The performance increase for Excel from SGI Altix XE clusters also means that extremely complex, iterative calculations, such as Monte Carlo simulations or parametric sweeps, which in the past typically required overnight batch processing, can now run during the day. WCCS's pro-active job monitoring and failover capabilities increase overall reliability for such types of applications.

WCCS on SGI Altix XE is also of key advantage to industries where users have long relied on third-party applications written for the Windows environment. One such industry is manufacturing, where numerous ISVs offer CAE tools. Many of these applications have been extended to take advantage of the parallel processing capabilities of WCCS. As businesses migrate to HPC, both developers and end users can continue to use the tools that they are familiar with. The list of CAE ISVs who have applications that run on WCCS is impressive. It includes Abaqus, Ansys, CD-Adapco, Dassault Systemes, ESI Group, Fluent, LSTC, The Mathworks, MSC Software, and Wolfram Research.

6.0 SGI and Microsoft

SGI (Nasdaq "SGIC") is a leader in high-performance computing, with 25 years of experience in solving the most demanding compute and data-intensive problems. SGI delivers a complete range of high-performance server and storage solutions along with industry-leading professional services and support, enabling its customers to overcome the challenges of complex data-intensive workflows and accelerate breakthrough discoveries, innovation, and information transformation. With offices worldwide, the company is headquartered in Sunnyvale, California, and can be found on the Web at www.sgi.com.

Founded in 1975, Microsoft (Nasdaq "MSFT") is the worldwide leader in software, services and solutions that help people and businesses realize their full potential. For more information about Windows Compute Cluster Server 2003, please visit: http://www.microsoft.com/hpc

To join the HPC Community, please visit http://www.windowshpc.net

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

4058 [11.05.2007] **J15337**