WHITE PAPER

SGI® UV™ 2000 System Reliability, Availability and Serviceability

Optimizing RAS for the World's Most Scalable Coherent Shared Memory Computing Platform

June, 2012

Abstract

SGI UV 2000 architecture provides single-system image (SSI) scalability up to 4096 threads (256 sockets) and 64TB of memory. Given the potential size of these systems, and the importance of the applications that run on them, reliability, availability and serviceability (RAS) features are vital to ensuring continuous operation. Even as SGI pushes the number of supported processor cores and coherent shared memory (CSM) to new heights, the company is bringing innovative RAS features to high-end Linux® deployments. Building on the advanced RAS features found in Intel® Xeon® processor E5-4600 product family, the company's RAS efforts leverage unique experience building the world's largest and most robust server systems. Together, SGI and Intel® are making focused investments to enhance system reliability for large shared memory environments.

TABLE OF CONTENTS

1.0 Introduction

SGI UV 2000 was designed to address computational and data access challenges faced by some of the most demanding scientific and commercial applications, and to overcome some of the limitations of traditional compute clusters of smaller, thin-node servers. A growing number of applications in both scientific and enterprise settings require the ability to operate on, and analyze, extremely large data sets, requiring large numbers of processors and very large amounts of memory. Many of these applications benefit enormously from being able to fit a large data set or database entirely into system memory. At the same time, as traditional compute clusters grow in size — both in terms of the number of compute nodes and the number of individual processors — communication processing can become a substantial portion of the computation time required. SGI addresses both of these issues with the UV 2000 system. In SGI UV 2000, the NUMAlink® 6 ASIC on each UV 2000 compute blade performs communication processing in hardware, essentially fusing a large number of processors into a high-capability cluster that can function as a single Linux system. Large-scale coherent shared memory (CSM) architecture can provide entirely new ways of solving a wide range of problems in science, engineering and business that are not available with clusters of smaller compute nodes, enabling breakthroughs in areas that include:

- Fraud detection/cyber security
- Bioinformatics
- Data analytics
- Real-time interaction with massive data sets
- Memory-resident databases
- Dataflow applications

The architecture and fundamental design parameters of SGI UV 2000 systems provide important advantages for large-scale, data-intensive technical and enterprise computing. UV 2000 provides the ability to scale CSM, compute power and I/O bandwidth from very small to very large sizes. The SGI UV 2000 system (Figure 1) provides a multi-cabinet solution with up to 512 sockets for Intel® Xeon® processor E5-4600 product family (4,096 processor threads) and up to 64TB of shared memory in four to eight racks of hardware, depending on configuration. For more modest needs, the platform can scale much smaller size as well: as small as 4 CPU sockets or 16 cores and only a few gigabytes of memory.

Large, monolithic CSM systems bring value in terms of the problems they can address, and in terms of ease of use. Delivering on that value requires that these systems be extremely dependable. SGI UV 2000 advances the company's more than 25 years of achievement in delivering many generations of high-performance servers that operate reliably for years or even decades. These achievements cut across all aspects of the product — from design, manufacturing and testing to deployment, operation and the service cycle.

Figure 1. The SGI UV 2000 system contains up to 256 Intel® Xeon® processor E5-4600 product family and up to 64TB of memory in four racks (single rack shown).

2.0 SGI UV 2000: Built-In Dependability

Given the level of integration required to build large CSM systems, SGI has a strong collaborative partnership with Intel®. A variety of hardware and software enhancements have been made to the SGI UV 2000 platform for Intel® Xeon® processors to provide the reliability required for a system architecture. To enhance reliability, the architecture includes extensive fault isolation, data path protection, and monitoring and debugging functions to help ensure data integrity and prevent disruptions. System software has also been enhanced to identify problematic nodes and memory and to remove them from the active pool of scheduled resources.

With its high-end scalability and advanced RAS support, the latest Intel® Xeon® processor E5-4600 product family offers a dramatic increase in capability and value. In addition to up to eight high-performance cores, 16 execution threads and 20MB of level 3 cache, each Intel® Xeon® processor E5-4600 product family provides technology that can protect data, increase availability and minimize planned downtime. By utilizing these advanced processors, SGI UV 2000 can yield the highest levels of scalability, availability and data integrity — at a fraction of the price of proprietary mainframe and RISC architectures. Beyond the base functionality offered by Intel® Xeon® processor E5-4600 product family, UV 2000 offers considerable advantages as summarized in Table 1.

System Element	RAS Feature
System	• Full data path integrity • Firmware provisioning • FRU failure analysis • Online diagnosis • Uptime management
Blade Interconnect	• Full data path integrity • Auto retry on failure detection • Parity check on all interconnect memory-mapped registers and internal memories
Processors	• Boot-time isolation
Memory	• DRAM failure analysis • Page migration • Boot time disable • Tiered failure containment
Power and Cooling	• Redundant, hot-swappable power supplies and cooling fans • Online fault detection and ACPI support

Table 1. SGI UV 2000 systems provide a wealth of RAS features.

2.1 Reliability

As both system and memory sizes increase, reliability becomes increasingly important. Experience has shown that memory errors are the most common type of error.

- Soft memory errors can occur when alpha particles or cosmic radiation strikes a component, causing a storage cell to change its state unexpectedly.
- Hard memory errors typically occur with the failure of a memory bit, link or device.

In designing UV 2000, SGI engineers paid detailed attention to circuit design rules, as well as executing comprehensive verification and testing, to help avoid both soft and hard memory errors. Components have been carefully screened and selected to meet the company's tough performance and reliability standards. During manufacturing, highly stringent quality control procedures are further enhanced with special QA testing of specific customer configurations to assure that the system as a whole will function reliably once deployed.

Reduced physical complexity can contribute directly to system reliability. By using the latest CMOS VLSI design, manufacturing and blade packaging, SGI has dramatically reduced the component count in UV systems. Fewer components means reduced complexity, fewer mechanical joints and interconnects, and much higher availability of the resulting system. Where necessary, redundant components are provided in UV 2000. For example, all fans and power supplies are hot-pluggable and N+1 redundant. Additionally, all disks and I/O cards have hot-swap capabilities.

Preservation of data integrity is paramount to the company's RAS efforts. Traditional error-detection mechanisms are extensively employed, including memory, cache, register and interconnect data path checks, as well as end-to-end ECC (error-correcting code) data checks. In addition, the UV 2000 system has features to prevent silent data corruption, including:

- Enhanced memory error correction to correct up to 8 bits
- The ability to detect multiple failing DRAMs on a memory module
- Error scrubbing to minimize cumulative error growth
- Extensive data integrity assurance tools

Reliability also extends from the components used in the system, as well as from the operating system and storage subsystems employed.

- **Linux reliability.** SGI UV 2000 systems run the Linux operating system, and Linux features are key to system reliability. In fact, SGI is a leading contributor to the Linux community in all aspects of shared memory computing on Linux, including RAS. Specific features to prevent silent data corruption include internal state and consistency checks, application abort/system halt capabilities and extensive regression testing.
- **NUMAlink 6 technology.** Given its pivotal and central role in the system, the NUMAlink 6 ASIC is key to system reliability. All NUMAlink and XIO™ connections use cyclical redundancy checking (CRC) error detection and retry on all messages. NUMAlink 6 furthers data protection with ECC or parity on all memory-mapped registers and internal memories, address path parity protection, and soft error correction-in-transit.
- **Reliable storage technology.** Reliable storage is vital to providing reliable system operation, and all SGI UV 2000 servers support a variety of reliable storage products that provide features to enhance reliability. The XVM volume manager can automatically provide mirroring of important data to help ensure that no data is lost in the event of a disk failure. Fully fault-tolerant RAID units are also available. All SCSI or Fibre Channel storage units have redundant power, cooling and controllers, and can continue to serve data in the event of a failure. Hot plugging of disk drives is supported with RAID and Fibre Channel vaults to allow the replacement of failed drives without having to take storage offline. The system supports warm plug-in of disks, using an administrative command to shut down the bus before removing or inserting a disk.

2.2 Availability

Availability is a measure of how well an overall system responds to failed components or other unforeseen operation issues. SGI has considerable experience deploying high-performance systems in demanding environments, strengthening capabilities in this area. In the case of UV 2000, enhanced availability results from the ability to detect problems or failed components and safely remove them from the system, while also effectively handling issues related to power and environmental conditions.

In the event of a problem, sophisticated UV 2000 software minimizes the impact on system use. Power-on diagnostics run to check for problems every time the system is started. Detected problem areas — including both CPUs and memory — can be de-configured from the system, allowing the system to start and continue to function. The UV 2000 Chassis Management Controllers (CMCs) support the capture of specific data that is needed for field-replaceable unit (FRU) analysis once components are de-configured.

During normal operation, environmental sensors monitor temperatures and voltages to identify problems and gracefully shut down systems before a crash can occur. NUMAlink 6 on UV also features full data and address path protection, as well as cable disconnect/reconnect and signal auto-retry. Uninterruptible Power Supply (UPS) solutions are available, with a single solution covering the UPS system, power monitoring software and support. SGI also offers an optional integrated water-cooled door that efficiently removes machine heat before it reaches the datacenter.

2.3 Serviceability

Serviceability relates to the ability of a system to be serviced while it continues to operate. Many enhancements have been made to shorten the mean time to repair (MTTR) on SGI UV 2000 systems. The UV blade form factor greatly simplifies serviceability, since components can be easily accessed and removed from a system without the need to remove bulkheads and cables. Independent power within a module or rack makes it possible for parts of the system to be shut down for maintenance while the remaining parts continue to function. Sophisticated online diagnostics keep watch over ongoing system operations. UV systems also integrate a service node that allows for quick diagnostics and repairs before a technician arrives on site. A variety of services options are also available, and are described later in this document.

3.0 Reliable System Components

While not all hardware faults are preventable, careful design and selection of components can help reduce the number and impact of errors that do occur, through automatic error detection and correction. For example, UV 2000 systems are designed to reliably detect all errors in system memory, directories and data paths. SGI utilizes ECC (Error-Correcting Code) on all system buses and memories to detect and correct single-bit errors, and detect double-bit errors. Effective environmental control also contributes to system reliability.

3.1 Memory Scaling

All SGI UV 2000 systems utilize the Intel® Xeon® processor E5-4600 product family. With support for high-end scalability and advanced RAS features, the latest Intel® Xeon® processor E5-4600 product family provides a dramatic increase in capability and value for systems such as UV 2000. The Intel® Xeon® processor E5-4600 product family provides silicon-level support for new advanced RAS features (Table 2). Key RAS benefits of these processors include:

- Robust data integrity. Data errors are prevented, detected, corrected and contained more comprehensively and effectively to preserve data integrity. If an uncorrectable error does occur, it is tagged and contained to help prevent propagation to other systems and applications.
- Improved system availability. Machine Check Architecture Recovery (MCA Recovery) enables OS-assisted system recovery from certain uncorrectable errors that might have brought down previous generation servers.
- Enhanced serviceability. Enhanced error logging and reporting enables predictive failure analysis to identify problematic components before they cause downtime or uncorrectable errors.

Benefits	Silicon Features
Protects Data: • Reduces circuit-level errors • Detects data errors across the system • Limits the impact of errors	• Parity checking and Error Correction Code (ECC) • Memory thermal throttling • Memory demand and patrol scrubbing • Corrupt data containment mode • Intel® QuickPath Interconnect (Intel® QPI) protocol protection via Cyclic Redundancy Checking (CRC): 8-bit or 16-bit rolling
Increases Availability: • Heals failing data connections • Supports redundancy and failover for key system components • Recovers from uncorrected data errors	• Machine Check Architecture Recovery (MCA Recovery) • Intel® Scalable Memory Interconnect (Intel® SMI) lane failover • Intel® SMI clock failover • Intel® SMI and Intel QPI packet retry • Intel® QPI clock failover • Intel® QPI self-healing • Single Device DRAM Correction (SDDC) plus random bit error recovery • Dynamic memory migration
Minimizes Planned Downtime: • Predict failures before they happen • Maintain partitions instead of systems • Proactively replace failing components	• Electronically isolated (static) partitioning • MCA error logging (CMCI) • CPU on-lining

Table 2. Intel® Xeon® processor E5-4600 product family provides a wealth of RAS functionality

3.2 Innovative NUMAlink 6 ASIC

Each SGI UV 2000 compute blade is equipped with two sockets for the Intel® Xeon® processor E5-4600 product family, with up to 512GB RAM per blade. The two processor sockets interface with the rest of the system through a NUMAlink 6 ASIC (Figure 2). Besides supplying the basic building block for the UV 2000 system architecture, these components include error detection and correction capabilities that benefit the reliability of the overall system. In addition to providing hardware-accelerated communication processing and connectivity, the NUMAlink 6 ASIC provides several innovations that contribute directly to increased RAS capabilities of SGI UV 2000 systems.

- NUMAlink 6 protocols and the NUMAlink 6 ASIC have been enhanced with additional error checking and retry capabilities to reduce transient communications errors by two orders of magnitude.
- By off-loading remote memory reads to the NUMAlink 6 ASIC, failures that would have caused processor hangs in previous generation systems can instead be retried or dealt with gracefully.
- The NUMAlink 6 ASIC provides safe mechanisms to communicate between nodes, even in the presence of node, memory or interconnect failures.
- CRC (Cyclic Redundancy Check) is provided on all NUMAlink 6 channels between UV 2000 blades.

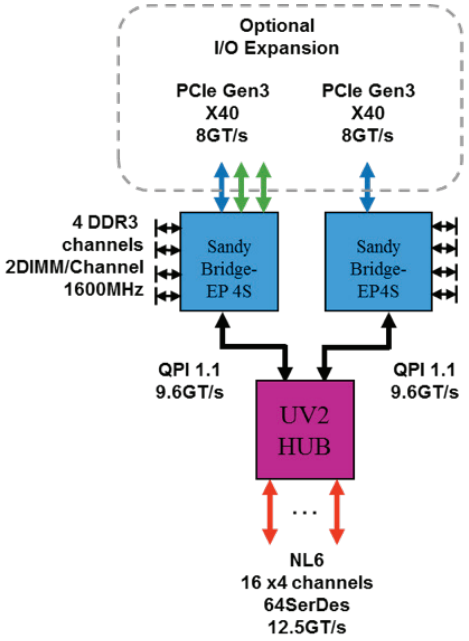

Figure 2. A NUMAlink 6 ASIC connects to two sockets for Intel® Xeon® processor E5-4600 product family on each SGI UV 2000 compute blade.

3.3 Environmental Monitoring, Power and Cooling

Environmental control is a critical element in system reliability. Clean power and proper cooling can dramatically improve observed reliability in a large system, avoiding overheated components that can rapidly become unstable and lead to more complex and unpredictable failures. SGI UV 2000 features an extensive environmental monitoring and control system to protect hardware operation:

- Redundant power and fans protect against failure in these components.
- Variable-speed fans ensure that the system always runs at the optimal temperature.

- Automatic system shutdown is provided in the (unlikely) case of over-temperature conditions, to prevent damage.
- Power efficiency contributes to cool operation and low failure rates.
- Power-efficient components help maximize performance while minimizing physical server footprint, power consumption and cooling requirements.

SGI UV 2000 systems feature a power supply and conversion architecture designed to convert AC power to DC voltages with over 90% efficiency (compared to efficiencies in the 60-70% range for many other vendors' products). Moreover, UV 2000 compute blades are designed to minimize power loss, with 12V DC blade input voltage requiring only one additional conversion to usable logic-level voltage. Finally, the Intel® Xeon® processor E5-4600 product family are some of the most efficient processors on the market, in some cases running at a fraction of the power of competitive RISC CPUs.

The superior energy efficiency of UV 2000 means that the vast majority of deployments do not require water cooling. However, SGI has a long history of deploying very large systems in densely-populated datacenter facilities, and was an early adopter of water cooling technologies for traditional air-cooled servers. Optional available water cooling for UV 2000 provides radiator-like cooling coils that intercept hot air as it exits each rack, efficiently cooling the air and preventing machine room hot-spots and the common problem of hot-aisle to cold-aisle recirculation. SGI water cooling systems can stabilize the ambient inlet air temperature and result in increased reliability.

4.0 Availability Features

SGI includes a number of features in UV 2000 systems that enhance overall availability beyond the reliability provided by the base hardware components. For example, as a leader in the development of high-end Linux solutions, SGI has made significant contributions to ensure the performance and reliability of Linux. SGI releases Linux enhancements that improve availability to the broader Linux community whenever they are generally applicable. Examples of SGI contributions relating to RAS functionality include:

- UCE (Uncorrectable Error) recovery enhancements
- Improved hardware error reporting
- Reduced panics on double-bit errors
- Better fault containment for cross-partition jobs

All such strategies represent potentially significant compromises in computational accuracy, performance and productivity.

4.1 System Partitioning

SGI UV 2000 hardware partitioning allows a single physical system to be subdivided into multiple logical systems without re-cabling. Partitioning capabilities are designed into the UV 2000 hardware to ensure ease-of-use along with highly reliable operation. In addition, while partitioning does change the OS configuration, memory used by applications can still be selectively shared across partitions using UV 2000 CSM capabilities. Through this technique, a UV 2000 system can be partitioned and run as a cluster for improved availability without losing all of the large-memory advantages of the architecture. UV 2000 hardware partitions can be rebooted independently without affecting operations in other partitions, providing a number of availability benefits:

- Necessary hardware repairs in one partition can be undertaken without disrupting other partitions.
- When upgrades are necessary, rolling kernel updates can be used to update each partition in turn without bringing the entire computing infrastructure to a halt.

- Organizations doing software development and testing can use partitioning to create development and test environments that closely approximate the production environment. These development and test partitions can be re-started as necessary — or brought down by ill-behaved software — without affecting other production partitions.

A number of unique hardware features increase the robustness of hardware partitioning on UV 2000 as compared with other systems.

- **Memory Protection.** UV 2000 has memory protection built into the SGI-designed chipset that resides in each UV 2000 compute blade. This feature provides fault containment by protecting each partition from unexpected writes from other partitions. Other systems that lack this hardware feature may be subject to memory corruption if a misconfigured kernel or poorly-behaved application attempts an inappropriate memory access. For example, XPMEM support in SGI MPI libraries allows the hardware to change memory protection on memory being shared with other partitions. As a result, Global Reference Units (GRUs) in one partition can directly load and store to shared memory without opening memory to access from all other partitions.
- **Reset Fences.** This capability is also built into the NUMAlink 6 ASIC to protect a partition from hardware resets occurring in another partition. Reset fences ensure that each partition operates independently and reliably in the face of restarts or hardware and software failures occurring in other partitions, and provides support for concurrent replacement of system modules.
- **Global Reference Unit (GRU).** The GRU built into SGI UV 2000 ASICs provides a reliable way to transfer data between partitions. This capability allows partitions to share data via high speed copying, if desired. To ensure fault containment, the GRU is designed so that a disruption in a remote partition will not crash or hang a partition that is actively performing a remote reference.

4.2 Memory Enhancements

Memory errors remain the most common errors experienced by servers. In typical Linux environments, memory configurations are relatively modest compared to those achievable with UV 2000 systems. For this reason, SGI is committed to improving the robustness of Linux for large memory systems.

When a memory location is determined to be bad because it has exceeded a threshold number of corrected ECC errors, the UV 2000 memory flawing feature allows the OS to move the data to a different page and mark the page containing that memory as flawed. The operating system subsequently avoids using the flawed page. SGI has also enhanced Linux to ensure that if an MCA event due to a hardware failure disrupts the system, the complete hardware state is captured. This ability improves root cause failure analysis to help ensure that the right components are quickly replaced to restore the system to full operation.

4.3 Reliable I/O

SGI has long been a leader in Fibre Channel and high-performance storage. In fact, SGI pioneered the deployment of a highly redundant Fibre Channel storage infrastructure along with the software necessary for efficient utilization. SGI UV 2000 profits directly from the reliability benefits of robust storage systems, with features that include:

- **Multi-path I/O.** Systems with multiple Fibre Channel host bus adapters spread across multiple blades and are connected either directly or through a fabric to SGI InfiniteStorage™ RAID arrays, resulting in I/O infrastructures with no single points of failure. Multi-path I/O balances I/O load across channels and shifts the load from a failed port or HBA over to survivors.

- **InfiniteStorage File System (XFS®).** Created by SGI, and now available in standard Linux distributions, XFS accommodates the I/O requirements of high-performance computing (HPC) environments while providing the reliability of journaling for error recovery and rapid restarts.
- **SGI InfiniteStorage Shared File System (CXFS™).** For shared data access in clusters, CXFS builds on XFS to create a highly reliable, high-performance storage infrastructure that lets cluster members read and write data directly to disk at full SAN speeds. UV 2000 systems that have been partitioned can use CXFS to allow partitions to share access to the same data sets without compromising performance.

4.4 High Availability Cluster Configurations

To further enhance UV 2000 system availability, clustered configurations can be designed using SGI InfiniteStorage Cluster Manager for Linux or other third-party high-availability software solutions for Linux. With Cluster Manager, highly-available application services can be created that span separate UV 2000 systems, or partitions in a single UV 2000 system. Applications fail over from one cluster member to another, should anything affect the running service.

5.0 Maximum Serviceability

The company's goal is to make most components of the SGI UV 2000 platform serviceable by an administrator with minimal or no system disruption. The UV 2000 system design includes important serviceability features that provide advanced system control capabilities, system health monitoring, online system management and maintenance, and failure analysis. By its very nature, the advanced, modular blade design of UV 2000 enhances serviceability, with individual system components easily accessed for service, maintenance or upgrade.

5.1 SGI UV 2000 Features to Maximize Uptime During Servicing

UV 2000 compute blades are housed in a chassis referred to as an enclosure. The UV 2000 enclosure contains up to eight compute blades. A single UV 2000 rack houses up to 4 enclosures (total of 32 blades, 64 sockets, 512 cores) as shown in Figure 3.

Figure 3. An SGI UV 2000 enclosure (four per rack) supports up to 8 UV 2000 compute blades.

Components in each enclosure are electronically isolated so that they can be replaced without powering down the enclosure. In most cases, power supplies and individual PCI cards can be hot-swapped without interrupting the operation of the system or partition containing the component. If a system has been partitioned, a failed compute blade in one partition can be replaced without affecting the operation of other partitions.

However, the partition containing the failed compute blade must be shut down while the operation is in progress. Other related features of the UV 2000 system help minimize downtime caused by component failure, including:

- The advanced RAS capabilities of the Intel® Xeon® processor E5-4600 product family minimize the likelihood of a CPU failure.
- Compute blades can be disabled and a system can run without them until scheduled maintenance becomes possible.
- Individual memory pages can be marked as flawed and retired while operations continue.
- Processors and memory are always subjected to self-test at boot time and automatically de-allocated if failures occur. The system is then able to boot without the affected resource so that operations can continue.

5.2 System Management Network

All UV 2000 enclosures contain an embedded Chassis Management Controller (CMC) that runs on standby power and is operational whenever the enclosure is connected to an active power source. The CMC network manages the hardware partitions within each system, providing pinpoint power control, system booting and support for configuration control. CMCs can transparently extract all internal register states and actions from the attached UV 2000 compute blades while the system is running, providing a wealth of input data that allows a fault analyzer to produce failure data reports down to the field replaceable unit (FRU) level.

The CMC is able to read the complete hardware configuration down to the level of individual FRU serial numbers in real-time. This capability supports the rapid and accurate notification and transmittal of essential information for system service actions. The CMC provides control and monitoring functionality for each enclosure, as well as communication to other CMCs. The CMC is active even when the system is not booted or powered off. Overall, the system controller network provides the following functionality:

- Power control to the entire system
- Power control to individual compute blades
- Environmental monitoring
- Monitoring status and error message information
- Specific commands to monitor or change system functions
- System boot control

Multiple blade chassis in a system are managed by a System Management Node, which is available in a redundant, failover high availability configuration to assure that system management functions are not lost in the unlikely event that the main management node fails.

6.0 SGI Service Options to Assure Uptime

As a long-standing global leader, the SGI Customer Support organization offers a broad range of services, up to and including 7x24 system support. The SGI Customer Support organization consistently ranks among the top in the industry according to third-party evaluation metrics. Additional available services include:

- SGI MAS (Managed Services) Console. The SGI Solution for Console Server Management is a valuable tool to help system administrators monitor and manage SGI servers during a system-down situation

by providing an interface to the system even when network access is not available. The solution is a combined hardware, software, support and on-site installation package for management of one or more heterogeneous servers.

- SGI UpSafe Uninterruptible Power Supply (UPS). UPS systems are critical to protect electronic equipment from power problems such as blackouts, brownouts and electrical surges/sags caused by the weather or by events such as the switching off of heavy industrial equipment (elevators, factory machines, etc.). A UPS is especially critical in systems located in a multi-tenant building where there is competing demand for power and momentary blackouts are common. Through UpSafe, SGI delivers a full line of UPS solutions configured to meet the specific needs of the datacenter environment and the UV 2000 server configuration. A single solution covers the uninterruptible power system, power monitoring software and support.
- SGI Electronic Support - Embedded Support Partner. SGI Electronic Support is a fully integrated suite of services that work together seamlessly to monitor and manage the system and proactively protect against problems. For example, if an anomaly is detected while SGI Embedded Support Partner (ESP) is activated, SGI Electronic Support notifies the system administrator and an SGI professional, searches for possible fixes, and then sends field-proven solutions related to the identified system problems. SGI Electronic Support provides a smooth, seamless, multipoint support experience that reduces downtime, minimizes administrative resources and saves money. In many cases, problems may be detected and corrected before users are even aware of them. SGI Electronic Support is available at no additional cost to system owners who have a valid SGI Warranty, SGI FullCare™, SGI FullExpress™ or SGI FullExpress™ 7x24 support contract.

7.0 Conclusion

The SGI UV 2000 product line represents a new level of performance, serviceability and overall capability that builds on the strong tradition of powerful and available shared memory systems from SGI. Along with support for up to 2,048 processor cores and up to 64TB of memory, SGI UV 2000 RAS capabilities demonstrate a trend of constant improvement. Based on a close and fruitful collaboration with Intel®, the SGI UV 2000 system combines the RAS features of the Intel® Xeon® processor E5-4600 product family with the innovative NUMAlink 6 interconnect to create the reliable, available and serviceable UV 2000 system architecture from SGI.

SGI considers the RAS capabilities of our flagship systems to be a vital and integral component of a reliable and dependable high performance computing solution. The RAS infrastructure of UV 2000 systems — including the system controllers, operating system software, diagnostics and internal firmware — undergo steady RAS capability improvements as techniques for monitoring and managing systems are refined and optimized in customer environments. SGI manufacturing also employs strong feedback and process control methods for continued enhancement of reliability. Future development will include continuous advancement in RAS capabilities based on customer feedback and real-world experience, enabling SGI to continue to deliver some of the largest high performance computing systems available.

Global Sales and Support: sgi.com/global