WHITE PAPER

SGI® UV™ 3000 System Reliability, Availability and Serviceability

Enabling continuous operation of high performance computing for compute-intensive workloads

August, 2015

Abstract

Designed for high performance compute-intensive, fast algorithm applications such as CAE, genome assembly, and scientific simulations, SGI UV 3000 is an advanced symmetric multiprocessing (SMP) system delivering up to 4096 cores (256 sockets) and 64TB of cache-coherent shared memory. Key to hosting applications at extreme scale, eliminating application silos through system consolidation, or running ever larger technical applications requiring many hours run time, the UV 3000 is also equipped with high fault-tolerance and robust reliability, availability and serviceability (RAS) features to help ensure continuous operations. SGI innovation leverages many years' experience delivering the most powerful computing systems in the world, and together with Intel®, SGI is making focused investments to further enhance system resiliency for technical applications.

TABLE OF CONTENTS

1.0 Introduction

SGI® UV™ 3000 was designed to address the computational and data access challenges facing large technical computing environments by overcoming the limitations associated with traditional High Performance Computing (HPC) clusters. A growing number of technical computing applications require the ability to process and analyze extremely large data sets, requiring large numbers of processors combined with very large amounts of memory and high performance IO. Many of these applications can benefit enormously if a large data set can be placed entirely into system memory versus residing on disk, or even FLASH storage. Distributing data across nodes in a server cluster, however, has performance limitations and introduces complexity. As 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 adding latency. Thus certain applications require placement on a “single node” so that data can be accessed and analyzed holistically. This presents a challenge for large data environments, resulting in application silos and the need for data sampling.

SGI addresses this challenge with the SGI UV 3000. Featuring the high bandwidth, ultra-low latency SGI NUMAlink6 ASIC, large numbers of processors are essentially fused to form a single high capacity, cache-coherent shared memory Linux system. This 6th generation coherent-shared-memory computing architecture provides entirely new ways of solving a wide range of problems in technical environments not possible with clusters of smaller compute nodes, enabling breakthroughs in areas spanning:

- Computer-Aided Engineering (CAE)
- Genome assembly
- Scientific simulations
- Preemptive fraud detection/cyber security
- Data streaming/filtering

The architecture and fundamental design parameters of the UV 3000 provide important advantages for large-scale, compute-intensive environments through its ability to scale cache-coherent shared memory, compute power, and I/O bandwidth from small to very large sizes. Shown in Figure 1, SGI UV 3000 scales from 4 to 256 sockets utilizing Intel® Xeon® E5-4600 v3 processors (up to 8,192 processor core-threads) and up to 64TB of shared memory.

A large, monolithic coherent-shared-memory system has high value in terms of the problems it can address, as well as ease of use. Delivering that value at high service levels also requires the system be extremely dependable. SGI UV 3000 advances the company’s quarter-century history of achievement in delivering high performance servers that operate reliably for years, even decades. These achievements cut across all aspects of the system — from design, manufacturing and testing, to deployment, operation and the service cycle.

Figure 1: The SGI UV 3000 system contains up to 64 Intel® Xeon® E5-series processors and up to 16TB of shared memory in one rack. Up to 4 racks can be combined in a single coherent system with up to 4,096 cores and 64TB of memory.

The SGI RAS philosophy is to minimize the potential for unplanned outages and in rare cases where a problem occurs, provide an accurate diagnosis and rapidly carry out a repair – either through pro-active management of resources while the system continues to operate at reduced capability, or through a service action – to minimize the forward impact. In addition, SGI approaches RAS in a holistic way by taking into account the need to deliver the highest levels of performance combined with RAS. It should not be an either/or proposition or purely a checklist of capabilities, as additional features designed to improve RAS can negatively impact performance. This can be the case if a given feature is overly complex and resource hungry in its implementation – and which can lead to an actual reduction in real reliability.

This paper outlines methods employed by SGI to deliver RAS for the SGI UV 3000 and enable continuous operation of high performance in-memory computing for compute-intensive workloads.

2.0 SGI UV 3000: Design Process

To deliver reliable, stable, scalable systems, SGI has built upon years of investment in meticulous engineering practices. These include design practices used for ASIC hardware, high speed interconnects, high-speed PCB design, platform software development, and careful component selection.

SGI designed the UV 3000 to utilize Intel® Xeon® E5 x86 processors, which offer the strong support for RAS features and bring the full weight of Intel® design resources to deliver a high quality component.

Development of the SGI NUMAlink6 ASIC, the heart of the UV 3000, entails extensive design simulation to include fault injection techniques to ensure absolute robustness. Design of the printed circuit board (PCB) employs exceptionally high standards with built in margins to the design rules, signal integrity rules, and component selection – well beyond the rules used for typical reference designs. This best practice also extends to product fabrication and quality assurance, where extensive inspection techniques such as X-ray inspection are employed to reach deep inside complex circuit boards. A strong feedback loop between engineering and manufacturing ensures consistent quality. In addition, rigorous shock and vibration testing standards are used to complete the overall integrity and robustness of the system design. The high level of system reliability measured by customers demonstrates the outcome of this holistic design approach.

SGI UV utilizes a progression of resiliency methods that contribute to overall system availability through a number of resilient actions, redundancy, and careful monitoring to pro-actively address potential issues before they become critical. The SGI drive to achieve world-beating system performance, maximize memory capacity, and maintain high availability guides the design in many ways. Memory resiliency is a key example and among several principle advantages of an SGI system. Common industry practice is to use memory mirroring and memory sparing. This approach, however, substantially increases power consumption and cost while reducing capacity– all contrary to the goals of the UV 3000. SGI uses a tiered approach.

DRAM full device ECC, retried transactions, background memory scrubbing, device-failure mechanisms, and auto-correction form the base layer of memory resilience. These mechanisms are built upon with continual and extensive DRAM failure analysis, of which a key component is SGI's Memlog utility. Corrected memory errors are logged and analyzed. If a DIMM page (4k bytes in size) is deemed defective, an attempt is made to transparently relocate data to a new page and retire the old page, enabling applications to continue running without interruption and reducing exposure and uncorrectable errors. Administrators are also alerted to failing DIMMs in need of replacement during planned maintenance windows.

The UV 3000 chassis contains 8 blades, for which each blade contains an SGI UV ASIC with integrated NUMAlink6 interconnects that enables growing the coherent shared memory system at extreme scale. Once again, a layered approach is used to ensure system resilience. This includes features at the lowest levels of the interconnect fabric, including link-level-retry and software driven policies that pro-actively detect potential link failures, steer traffic away from a potentially weak link in the system, and schedule a planned service action to address root cause.

On top of a progression of mechanisms built into the base system infrastructure, SGI UV 3000 has the additional capability of running in High Availability (HA) mode using advanced HA failover technology, so workload service levels can rapidly transition from one system to another in the event of catastrophic failure, or during scheduled maintenance.

3.0 SGI UV 3000: Built-In Dependability

Given the level of silicon integration required to build large cache-coherent shared memory systems, SGI has a strong collaborative partnership with Intel®. A variety of hardware and software enhancements have been made to the SGI UV 3000 platform that leverage Intel® Xeon® processors to provide the reliability required by the 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 been enhanced to identify problematic blades 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-family offers a dramatic increase in capability and value. Each Intel® Xeon® E5-family processor provides technology that can protect data, increase availability, and minimize planned downtime. Using these advanced processors, the SGI UV 3000 leverages the highest levels of scalability, availability and data integrity at a fraction of the price of proprietary mainframe and RISC architecture. Key Intel® Xeon® E5-family RAS benefits include:

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): • Checksum of 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® 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)

Table 1: Intel® Xeon® E5-series processors provide a wealth of RAS functionality

In addition to the base functionality offered by Intel® Xeon® E5-family processors, UV 3000 offers considerable additional RAS features as summarized in Table 2.

System Element	RAS Feature
System	• Full data path integrity • Firmware provisioning • FRU failure analysis • Online diagnosis • Uptime management
NUMALink6 Interconnect	• Full data path integrity • Auto retry on link level failure detection • Alpha immune latches used in ASIC design • Predictive routing around potential faults • Adaptive tuning and re-tuning • Link quality monitoring
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 2: SGI UV 3000 systems are designed with numerous RAS features

3.1 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 transistor to change its state unexpectedly.
- Hard memory errors typically occur with the failure of a memory bit, link or device.

In designing the UV 3000, SGI engineers optimized performance and reliability of memory. Particular attention was applied to circuit design rules, while comprehensive verification and testing techniques were applied (building upon previous generations of system design experience) to ensure reliable operation, to help avoid both soft and hard memory errors. Component suppliers have been selected to meet the company's tough performance and reliability standards. During manufacturing, quality control procedures are further enhanced with specialized testing of specific customer configurations to ensure that the system as a whole will function reliably once deployed.

Reduced physical complexity contributes directly to system reliability. By using the latest CMOS VLSI design, manufacturing, and advanced 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. Redundant components are provided where necessary in UV 3000. For example, all fans and power supplies are hot-pluggable and at least N+1 redundant (some configurations can be N+N redundant on power supply, with full A/B power configuration redundancy). Additionally, all disks 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 3000 has features to prevent silent data corruption including:

- Enhanced memory error correction to correct up to 4 bits
- The ability to detect multiple failing DRAMs on a memory module
- Error scrubbing to minimize cumulative error growth
- Extensive data integrity assurance tools
- Hardware error injection capabilities to verify proper error detection and recovery

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

- **Linux reliability.** UV 3000 systems run the Linux operating system and Linux features are key to system reliability. 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.
- **NUMAlink6 technology.** Given its central role in the system, the NUMALink6 ASIC is key to system reliability. All NUMAlink connections use cyclical redundancy checking (CRC) error detection and retry on all messages. NUMALink6 also provides data protection with ECC or parity on all memory-mapped registers and internal memories, address path parity protection, soft error correction-in-transit and alpha immune latches.
- **Reliable storage technology.** Reliable storage is vital to providing reliable system operation and all SGI UV 3000 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 SGI 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.

3.2 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. For UV 3000, 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 3000 software minimizes the impact on system use. Power-on diagnostics 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 3000 Rack Management Controller (RMC) captures specific data needed for field-replaceable unit (FRU) analysis once components are de-configured.

During normal operation, environmental sensors monitor temperatures and voltages to identify stress areas and gracefully shut down systems before a crash can occur. NUMALink6 also features full data and address path protection, as well as cable disconnect/reconnect and signal auto-retry. NUMALink6 can also reroute interconnections around failing links, preventing system crashes. 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. A more detailed look at NUMALink6 and environmental monitoring of power and cooling follow in this document.

3.3 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 3000 systems. Independent power within a chassis 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 3000 systems make use of a rack management controller that allows for quick diagnostics and repairs before a technician arrives on site. A variety of services options are available and described later in this document.

4.0 Innovative NUMALink6 ASIC

Each SGI UV 3000 chassis is equipped with eight 2-socket blades using Intel® Xeon® E5-series processors and with up to 16TB of RAM using 64GB DIMMs (up to 512GB per socket). The two processor sockets per blade interface with sockets in other UV 3000 chassis through a NUMALink6 ASIC (Figure 2). Serving as the heart of the UV 3000 system architecture, these ASICs 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 6th generation ASIC has several innovations that contribute directly to increased RAS capabilities of UV 3000 systems:

- NUMALink6 protocols and the NUMALink6 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 NUMALink6 ASIC, failures that would have caused processor hangs in previous generation systems can instead be retried or dealt with gracefully.
- The NUMALink6 ASIC provides safe mechanisms to communicate between nodes, even in the presence of certain node, memory or interconnect failures.
- CRC (Cyclic Redundancy Check) is provided on all NUMALink6 channels between UV 3000 blades and UV3000 chassis.
- NUMALink6 can often reroute paths around failing links, maintaining connectivity and system stability.

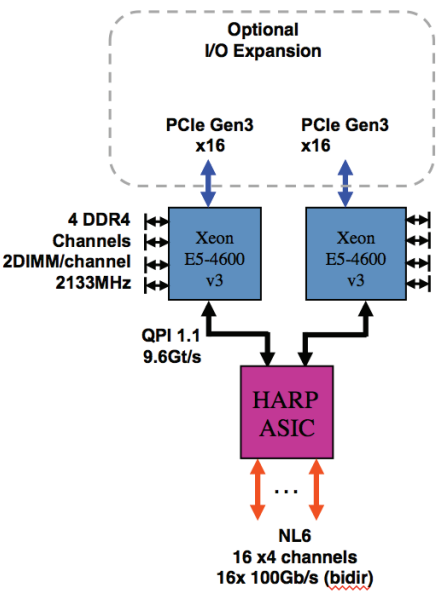

Figure 2. Each NUMALink6 ASIC connects to 2 Intel® Xeon® processor sockets on each SGI UV 3000 blade.

4.1 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 3000 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.

UV 3000 systems feature a power supply and conversion architecture designed to perform AC/DC conversion with over 90% efficiency (compared to efficiencies in the 60-70% range for most other vendors). Moreover, UV 3000 chassis are designed to minimize power loss - the 12V DC chassis input voltage requires only one additional conversion to usable logic-level voltage. Furthermore, the Intel® Xeon® E5-series processors are some of the most efficient processors on the market, and in some cases, run at a fraction of the power of competitive RISC CPUs.

The superior energy efficiency of UV 3000 means that the vast majority of deployments do not require water-cooling. SGI has a long history of deploying very large systems in densely populated datacenter facilities, however, and was an early innovator of water-cooling technologies for traditional air-cooled servers. Optional available water cooling for UV 3000 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.

System power supplies and fans are also monitored for failure or degradation, with alert notification to the administrator or remote services.

5.0 Availability Features

SGI includes a number of features in UV 3000 systems that enhance overall availability beyond the reliability provided by the base hardware components, beginning with Linux. As a leader in the development of high-end Linux solutions, SGI has made significant contributions to ensure the performance and reliability of Linux and releases enhancements to the broader Linux community when they are generally applicable. Examples of SGI contributions relating to RAS functionality include:

- UCE (Uncorrectable Error) recovery enhancements
- Improved hardware error reporting
- Migrate data away from pages with corrected memory errors
- Better fault containment for cross-partition jobs

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

The SGI Hardware Event Tracker (HET) suite included in SGI Foundation Software monitors system sensors and environmental conditions such as power supplies, cooling equipment, temperature, and voltages. HET provides understandable and actionable alerts for UV 3000. HET support requires running SGI Management Center 3.0. To improve performance, accuracy, and achieve high scaling HET relies on event notification rather than polling, including:

- Hardware Event Tracking configures alerts for hardware component sensors with pre-defined critical thresholds and notifies these events via SNMP and email (e.g. shutting down the system if a component overheats)
- Hardware Event Tracking manages defective sensors which may be intermittently sending false event notifications
- Hardware Event Tracking automatically recognizes the addition of new hardware components and configures the alerts accordingly
- Hardware Event Tracking configures the system management controllers to perform system shutdown if the room ambient air temperature reaches critical temperature.
- Additionally, PCIe Advanced Error Reporting (AER) is supported, providing the information necessary to quickly fix I/O problems.

5.1 System Partitioning

SGI UV 3000 hardware partitioning allows a single physical system to be subdivided into multiple logical systems without re-cabling. Partitioning capabilities are designed into the UV 3000 hardware to ensure ease-of-use along with highly reliable operation. UV 3000 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.

Figure 3. Examples of partitioning SGI UV 3000.

A number of unique hardware features increase the robustness of hardware partitioning on UV 3000 as compared with other systems.

- **Memory Protection.** UV 3000 has memory protection built into the SGI-designed chipset that resides on each UV 3000 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 NUMALink6 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 NUMALink6 ASICs provides a reliable way to transfer data between partitions, as well as share data between partitions if desired via high speed copying. 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.
- **Topology.** UV 3000 blades and chassis have multiple interconnect links NUMALink6 (HARP) ASICs to maximize performance and provide alternative paths for interconnect traffic in the case of any increase in CRC retries on NL cables, that indicate early signs of a failing connection. This then allows dynamic re-programing of the NUMAlink routes to avoid a suspect link.

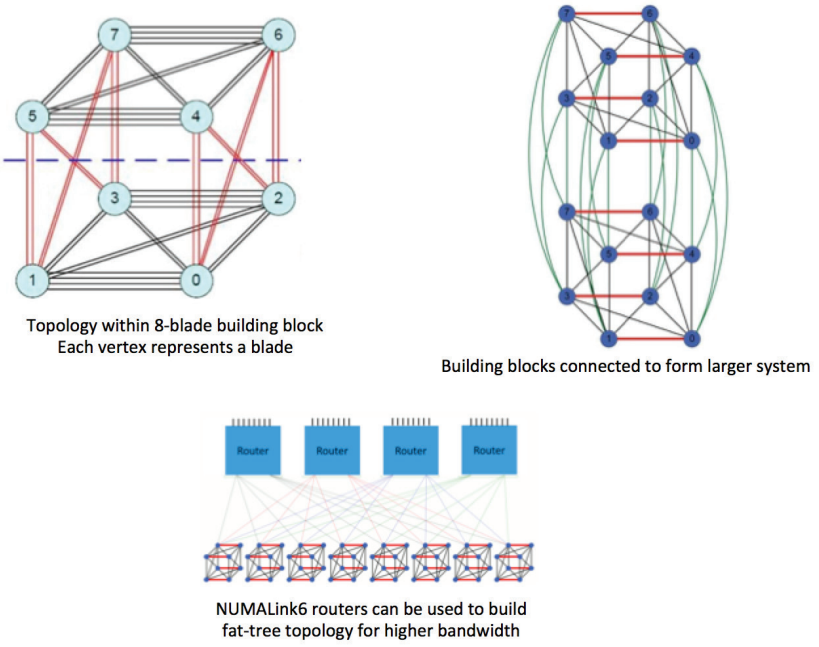

Figure 4. Topology options for SGI UV 3000.

5.2 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 3000 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 correctable errors, the UV 3000 memory-flawing feature in SGI Memlog directs 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.

5.3 Reliable Storage

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. UV 3000 profits directly from the reliability benefits of robust storage systems, with features that include:

- **Multi-path and Multi-Host I/O.** Systems with multiple Fibre Channel or SAS host bus adapters spread across multiple chassis 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. Multi-host enables a single volume to be shared across multiple hosts (required by CXFS) and used in some HA configurations.
- **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 3000 systems that have been partitioned can use CXFS to allow partitions to share data sets without compromising performance.

5.4 Linux High Availability

To deliver uninterrupted user access to mission-critical applications and leveraging hardware and software redundancy, UV 3000 systems can be configured in a Linux HA cluster with independent nodes each running separate copies of the Linux operating system and application software.

UV 3000 HA clusters ensure availability through the failover of applications and resources that the applications require such as storage, IP addresses, and software modules. If one server fails another server in the cluster, either local or remote, can assume the workload and deliver increased application uptime.

6.0 Maximum Serviceability

The UV 3000 includes important serviceability features that provide advanced system control capabilities, system health monitoring, online system management and maintenance, and failure analysis. SGI's goal is to make most components of the UV 3000 serviceable by an administrator with minimal-to-no system disruption. In addition, the innovative chassis design allows individual system components to be easily accessed for service, maintenance or upgrade.

6.1 SGI UV 3000 Features Maximize Uptime During Service

UV 3000 chassis contains 8 blades, with each blade hosting 2 processor sockets, up to 16 memory DIMMs (8 per processor), and up to 2 PCIe cards as show in Figure 3. A single rack fully populated with 4 UV 3000 chassis provides for a total of 64 processors and 32TB of shared memory.

Figure 5. An SGI UV 3000 chassis enclosure (up to 4 per rack)

The UV 3000 has hot-swap power supplies and fans, allowing replacement without interrupting system operation. If a system has been partitioned, other failed hardware components within one partition can be replaced without affecting the operation of other partitions. The partition containing the failed component must be shut down, however, while the service operation is in progress. Other related features of the UV 3000 help minimize downtime caused by component failure include:

- The advanced RAS capabilities of the Intel® Xeon® E5-series processors minimize the likelihood of a CPU failure.
- UV 3000 chassis can be individually disabled until scheduled maintenance becomes possible while remaining chassis in the system continue to run.
- Individual memory pages can be marked as flawed and retired without interrupting operations.
- Processors and memory are self-tested at boot time and automatically de-allocated if failures occur.
- The system is able to boot without the affected resource so that operations can continue.

6.2 System Management Network

All UV 3000 systems feature a Rack Management Controller (RMC) to provide management services. The RMC has a dedicated Ethernet connection to the Chassis Management Controller (CMC) that is provided in each UV 3000 chassis. In turn, each CMC has a dedicated Ethernet connection to the Blade Management Controller (BMC) on each of the 8 blades in a chassis. The RMC can transparently extract all internal register states and actions from the attached UV 3000 chassis 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 RMC is the control point for managing hardware partitions within each system, providing pinpoint power control, system booting and support for configuration control.

The RMC 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 RMC provides control and monitoring functionality for each UV 3000 chassis. The RMC is active even when the system is not booted or powered off. Overall, the RMC provides the following functionality:

- Power control to the entire system
- Power control to individual UV 3000 chassis and blades
- Environmental monitoring
- Monitoring status and error message information
- Specific commands to monitor or change system functions
- System boot control
- System health monitoring

7.0 SGI Service Options to Assure Uptime

The SGI Customer Support organization consistently ranks among the best in the industry according to third-party evaluation metrics and offers a broad range of services, up to and including 7x24 system support. Available services include:

Professional Services
SGI consultants average over 20 years of experience and bring broad industry and technical knowhow to provide tailored solutions to complex computing problems.

Installation Services
A broad range of product-focused services precisely defined and delivered by expert SGI engineers, help optimize system utilization and accelerate time to value.

Support Services
Personalized and quality service delivered via support programs that match your business needs, providing flexible coverage hours, priority response, and with no per-incident charges.

Customer Education
To fully maximize your SGI investments, we provide system administrators, system maintainers, and programmers with valuable knowledge about our systems and help them gain experience.

SGI Remote Services
To help ensure business continuance, SGI Remote Services provides 24x7 system monitoring and change or fault notification to SGI customer support. A secure connection capability enables customers to authorize SGI support personnel to diagnose issues real time, with key configuration and diagnostic information available on demand to shrink time to resolution.

8.0 Conclusion

Building upon a long-standing tradition of in-memory computing leadership, SGI UV 3000 delivers new levels of performance, simplicity, and fault-tolerance. With up to 4,096 processor cores and 64TB of cache-coherent shared memory, UV 3000 combines the RAS features of Intel® Xeon® E5-series processors and Linux with extensive SGI RAS technology to achieve high service levels for demanding technical computing environments.

SGI considers RAS capabilities a vital and integral component of high performance computing solutions. SGI RAS infrastructure — including system controllers, operating system software, diagnostics and internal firmware — undergo steady capability improvements as techniques for monitoring and managing systems are refined and optimized in real-world customer environments. This enables SGI to consistently deliver the most powerful – and dependable – high performance computing systems in world.

9.0 About SGI

SGI is a global leader in high performance solutions for compute, data analytics and data management that enable customers to accelerate time to discovery, innovation, and profitability. Visit sgi.com for more information.

Global Sales and Support: sgi.com