WHITE PAPER

Optimized Application Migration for SGI Single System Image (SSI) Architectures

Migrating Linux Applications from Intel® Itanium®-based SGI® Altix® to Intel® Xeon®-based Altix® UV Systems

1. Introduction

This white paper covers the process of migrating Linux applications from existing Intel Itanium-based SGI Altix systems to the new Intel Xeon-based Altix UV.

While Itanium-based systems will continue to be offered in the Altix product family, many customers are choosing to migrate to the new Altix UV platform for a variety of technical, performance and economic reasons as well as the flexibility of the wide variety of applications supported on the Intel Xeon-based platform.

Altix UV is based on a blade architecture tied together by the NUMAlink® 5 fabric. Each blade allows the system to scale by two sockets (8-16 cores) at a time in a Single System Image (SSI), producing a system with a globally-addressable memory space. The current limit within a unique SSI is 256 sockets (2048 cores) and 16TB of shared memory.

Migrating an application from one architecture to another can vary from being a challenging and extensive project to being only a relatively minor implementation detail. Managing the differences between the processor architectures, data and object code formats, compilers, and supported devices on the system are just some of the areas to address. For instance, application migration from MIPS-processor-based Origin® systems to the Intel Itanium-based Altix systems was challenging due to the differences between the operating systems (UNIX to Linux), and the processor differences (32 and 64-bit MIPS processors versus 64-bit Itanium processors). On the other hand, application migration from Itanium-based Altix systems to Xeon-based Altix UV can be relatively minor in most cases since applications typically are compiled with the same software development tools.

The Intel Xeon processors used in Altix UV use the Intel 64 architecture (commonly referred to as x86-64). This is a 64-bit superset of the x86 instruction set that has been in use for years by most 32-bit Intel processors. Intel's codename for the processor used in Altix UV is "Nehalem-EX", a member of the same processor family that SGI already uses in the Altix XE, Altix ICE, CloudRack™ and Rackable™ product lines. Because of the wide use of Xeon processors, both SGI and its customers are already familiar with the process of porting applications to this popular processor.

As a result, one of the compelling advantages of Altix UV is that application migration to the platform is relatively simple for customers whose applications are developed for x86-64 and ia64. However, some considerations do need to be taken into account. After a backgrounder on the Altix UV development environment (section 2), this paper discusses the most important porting considerations for both in-house and ISV applications.

2. The Altix UV Development Environment

2.1 What Should Application Developers Expect?

In most cases, developers simply recompile their existing Itanium-based applications to support Xeon processors and then execute them as before on Altix UV. Some specific guidelines for developing in-house applications and for commercial applications for SGI hardware are provided in section 3.

Similarly, the Linux libraries and tools customers are familiar with, including SGI ProPack™, are available for the Altix UV system. In addition, Intel's compilers, libraries, and software development tools for Linux support ia32, x86-64, and ia64 architectures. This further strengthens software development consistency between processor architectures.

2.2 Compilers and Libraries

Altix and Altix UV systems both support the same compiler tools (Intel compiler suites and GNU compilers). In fact, the compilers are compatible with both ia-64 and x86-64 architectures. Developers simply have to specify the correct flag to produce binaries for their target architecture. Likewise, both SGI and Intel support versions of the same application libraries for either architecture.

2.3 A Rich System Software Environment

SGI offers a comprehensive set of system software that speeds time to productivity on the Altix UV platform, mirroring what is available on the Itanium-based Altix system (see Figure 1, below).

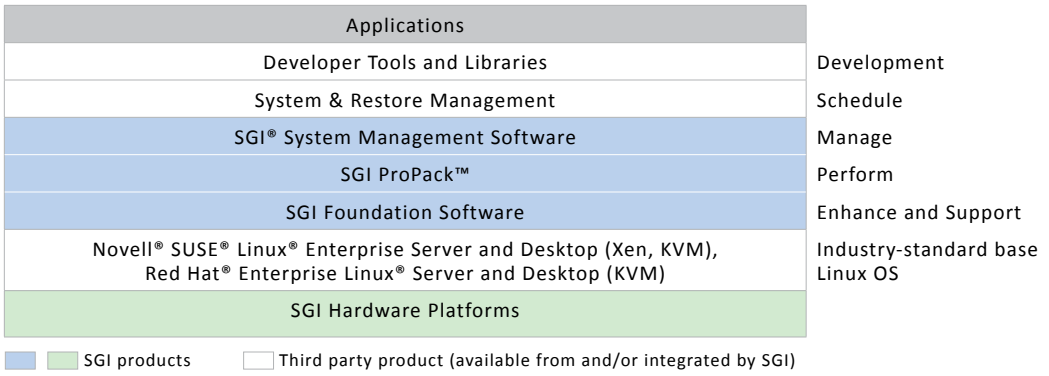

Figure 1. SGI's rich Linux software stack will be available on the Altix UV for consistency with existing platforms.

To ensure that Altix UV users obtain the best possible performance out of their systems, SGI ProPack is available with every Altix UV system. SGI ProPack 7 runs on Novell SUSE Linux Enterprise Linux 11 or Red Hat Enterprise Linux and provides functionality in four areas: accelerating applications, enabling parallel programming, supporting real-time performance and analyzing system resources (see Figure 2, below).

Accelerate Applications	**Parallel** Programming
• Cpusets optimize cpu and memory utilization • NumaTools data placement tools • Linkless Flexible File 1/O (FFIO) improves runtime I/O performance • SGI Fast Sparse Solvers • SGI MPI PerfBoost accelerates HP-MPI, Intel MPI, Open MPI applications	• SGI Message Passing Toolkit (MPI) • Perfcatcher MPI profiling tool • Easy launch of MPI Apps with Array Services • SHMEM message passing API • XPMEM cross partition support • UPC Compiler for UV systems
Real-Time Performance	**Manage** System Resources
• SGI REACT for Linux hard real-time performance • SGI Linux Trace debugger for real-time apps	• Performance Co-Pilot system performance monitoring • Comprehensive Systems Accounting • XVM volume management tool to stripe a volume across multiple physical disks

Figure 2. SGI ProPack 7 features.

3. Porting from Itanium-based Altix to Altix UV

3.1 Altix UV System Architecture

The system architecture of Altix UV is very similar to the existing Itanium-based Altix systems. Most importantly, Altix UV has a similar NUMAflex architecture with global shared memory, enabling a large number of processors to be supported in a single system image and the ability to share memory across multiple system images.

At launch, Altix UV scales to 2048 cores and 16TB of memory under a Single System Image. The architecture is designed to eventually scale to 16,384 blades with 32,768 sockets providing 262,144 cores (with eight cores per socket) for an overall performance greater then two petaflops!

3.2 Taking Advantage of Altix UV's MPI Offload Technology

SGI has made a number of significant architectural enhancements based on its long history in the High Performance Computing (HPC) market. The largest of these is the addition of an MPI Offload Engine (MOE) to the UV-hub ASIC. The MOE is similar in concept to a TCP/IP Offload Engine (TOE), which offloads TCP/IP protocol processing from system CPUs. Offloading MPI communications processing to the UV-hub accelerates critical

MPI communication while reducing CPU overhead and memory access latency, allowing MPI applications to scale to greater numbers of processors. SGI's MPI library, called the SGI Message Passing Toolkit (SGI MPT), takes advantage of the MPI Offload Engine without requiring any changes to the application.

Developing applications using SGI MPT is not the only way to take advantage of the MOE's performance optimizations. SGI also provides a tool called SGI MPI PerfBoost, included in SGI ProPack, which enables access to the MOE without the need for a recompile. MPI PerfBoost enables applications compiled with Platform MPI (formerly HP-MPI), Open-MPI and Intel MPI to be accelerated via the MOE. MPI PerfBoost enables the MPI calls from the application to use the equivalent SGI MPT call which is performance optimized by the MOE.

With the exception of MOE, most of the architectural differences in Altix UV derive from the change in processor architecture. Existing x86-64 Linux application binaries will run as-is on SGI Altix UV. Any Itanium-based Altix software directly related to supporting hardware devices (e.g. device drivers) will need to be ported to support Altix UV if one doesn't already exist.

3.3 Issues for In-house Developers

For customers who have created in-house applications on Itanium-based Altix systems, the porting process should be straight forward as the primary change is the processor architecture. SGI has successfully assisted with the porting of a large number of customer codes from ia-64 to x86-64 to support Altix ICE and its other x86-based systems. In general, the port to Altix UV will be even easier as its system architecture is closer to the legacy Altix architectures than other x86-based systems.

Even so, a few considerations should be taken into account:

- It is possible that some customer codes leverage ia-64 machine language to optimize certain functions. In such cases, the compiler would detect and report the problem, but it would be necessary for the customer to rewrite the code using x86-64 machine language to achieve the same result. Such occurrences should be rare.
- The next-generation Intel Xeon processor in Altix UV features an integrated memory controller. This means that server blades with two processors have two memory controllers, each tied to specific memory slots. As a result, processor A will have a slightly greater latency when accessing the memory of processor B and vice versa. While existing code will run without changes, advanced users may want to make optimizations in how their application accesses memory in order to achieve the highest possible performance.

3.4 Issues for Independent Software Developers

The porting issues faced by ISVs are expected to be fewer than normal since most ISVs will already have versions of their software compiled for Xeon-based systems. Only the subset of ISVs that have developed applications optimized to take advantage of unique features of the Itanium-based Altix platform will need to recompile. Some of these applications may be subject to the same considerations as discussed in the previous section (inclusion of machine code, NUMA effects within a node, etc.)

4. Conclusion

Developers migrating Itanium-based Altix applications to Altix UV systems should experience minimal problems or operational disruptions. For most, a simple recompile is all that is required. The same compilers, libraries and tools are available for both environments. A few developers may find that some additional work is required to achieve the same level of optimization on the new architecture.

If you would like more information about Altix UV, please visit the SGI website at www.sgi.com/altixUV or contact your local SGI sales office or channel partner.

Corporate Office
46600 Landing Parkway
Fremont, CA 94538
tel 510.933.8300
fax 408.321.0293
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 +61 2.9448.1463

© 2009 Silicon Graphics International Corp. All rights reserved. Altix, NUMAflex, NUMAlink, ProPack, REACT and SHMEM are trademarks or registered trademarks of SGI, Inc. Intel, Itanium and Xeon are trademarks or registered trademarks of Intel Corporation. All other trademarks, registered trademarks, tradenames, company names and service marks are the respective properties of their holders. 11112009 P/N 4194