SMP 'Workhorse' Enables New Avenues of Cutting-edge Research

SGI® UV™ 2000 Extends Compute Capabilities Beyond the Limits of Traditional Fat-node Clusters

Key Facts

Organization:
Masaryk University

Location:
Czech Republic

Industry:
Research

Business Overview

Masaryk University (www.muni.cz) is the second-largest public university in the Czech Republic. At the time of publication, it comprised nine faculties with more than 200 departments, institutes and clinics. Masaryk University is home to over 50,000 students, faculty and staff.

Recognized as one of the Czech Republic's most important teaching and research institutions, Masaryk University is committed to high quality education and research. In 2005 it received the European University Information System (EUNIS) Elite Award in recognition of its use of IT within the academic community – the first university in the Czech Republic to win this award. Together with Brno Technical University and several Czech Academy of Science Institutes it has also established The Center for Education Research, and Innovation in ICT (CEITEC) (www.ceitec.cz), the biggest research center in former Eastern Europe focused on leading edge life sciences and material science research.

Challenges

A vital resource for Masaryk University is the Institute of Computer Science, and CERIT Scientific Cloud (CERIT-SC) (www.cerit-sc.cz), its center for advanced computation, computer simulation and data processing. CERIT-SC's mission is to provide compute and storage resources and related services, including support for their experimental use.

Historically, CERIT-SC's computational resources have been provided by clusters made up of dual-socket servers, typically with 8-16 cores and 96-256GB memory per node. Currently, about 25% of these resources are formed by fat-node servers with up to 80 cores and 512GB memory, and two extra-fat nodes with 1.25TB memory. All of the clusters use an InfiniBand network to connect to each other, and to a common storage infrastructure.

Even with state-of-the-art technologies, CERIT-SC was finding it impossible to run particularly demanding tasks on its available infrastructure. This was the case in areas such as quantum chemistry (e.g.,, for tasks with more than 500 atoms), biologically active protein studies (e.g., cell membranes), molecular dynamics studies (e.g., whole virus simulation), and bioinformatics (e.g., error correction in genome sequencing). CERIT-SC anticipated similar problems in astrophysics (e.g., star formation), medicine (e.g., cell simulations), and other codes used within the Czech academic community, where there was growing demand but no adequate compute resource in the country where scientists could run such experiments.

To address these issues, CERIT-SC decided to invest in an Symmetric Multi-Processors (SMP) solution that was more powerful than that provided by its fat-node servers. They wanted a system with the memory capacity to handle applications greater than 4TB, a larger number of parallel processor cores to satisfy user demands, and the ability to run applications that are unsuited to cluster environments.

SGI Solution

Experts at CERIT-SC began to look for a system that would satisfy its most demanding users, and allow them to develop, study and test new applications without the need to be written for a distributed/cluster environment. Those environments dramatically increased the complexity of programming before it was even clear if their algorithm was usable.

sgi

A key question for CERIT-SC was whether it needed a hardware-based SMP implementation, or could achieve its goals through SMP software emulation running on a standard InfiniBand cluster. The team evaluated and tested a variety of possible SMP alternatives with at least 6TB memory. They found that software emulation could create many bottlenecks, and that some applications may become very slow. To harness the theoretical power of software emulation, substantial expertise and effort is needed. Which makes SMP implementation environments difficult to use as the distributed system of many cluster nodes – the environment CERIT-SC wanted to complement.

CERIT-SC's experts concluded that only an SMP server totally implemented in hardware would provide the performance and scalability for all the different workloads and programming techniques it required. And the only system capable of meeting its requirements was SGI® UV™ 2000, a solution whose capabilities had already been proven in many installations worldwide, and which passed all the stress tests specified by CERIT-SC to ensure its usability for a broad set of academic codes.

Further advantages of SGI UV 2000 were the availability of special hardware to accelerate typical constructs, such as MPI barriers and MPI reduce, which take a lot of time in typical cluster programming implementations; and the fact that throughput between SGI UV 2000 and external infrastructure could easily be scaled and extended in the future. After detailed comparisons of all possible configurations, CERIT-SC decided to complete its compute portfolio with the implementation of SGI UV 2000 with 48 Intel® Xeon® processor E5-4617 (288 cores), 6TB shared memory, and QDR InfiniBand and SAS for storage connectivity.

Results

SGI UV 2000 became CERIT-CS's SMP computational workhorse. Priority is given to applications and experiments requiring more memory and/or processing power than is available on CERIT-SC's fat-node servers, and the system is also used for in-house research and development work.

SGI UV 2000 is connected to a broad storage infrastructure, spanning several petabytes of storage capacity with different speeds and features, through six InfiniBand QDR lines (peak 3GB/s each) configured for High Availability. The system is also equipped with two sets of redundant disks, and attached via SAS to 70TB RAID for fast scratch usage.

Deployment of the system was completed with a three day training session attended by key users, CERIT-SC programmers and application specialists. The training focused on effective usage of modern processors and processor architectures – including advantages of the accelerators available in SGI UV 2000 – to help speed up the algorithms being developed by these users.

Benefits

The system was operational within a few days of delivery, since then it has been operating 24/7. Its reliability has been high, so even though there are nearly a thousand memory DIMMs and other components, failures do not interrupt system usage. Software stability has been very good; and the SUSE Linux Enterprise Server (SLES) operating system, SGI accelerators and tools have all worked as expected.

The good reputation and experience with the system also prompted researchers from CERIT-SC's sister organization CEITEC to purchase SGI UV 2000 to provide the resources and large shared memory required for its computationally intensive biochemical projects. This second system became operational in January 2014.

"SGI UV 2000 has delivered CERIT-SC an extraordinary performance for all tested codes, and its real workload performance even exceeded expectations of our experts The system also shows high long-term reliability and stability."

Aleš Křenek
Leader of HPC infrastructure at CERIT-SC

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.

For More Information

Please contact an SGI sales representative at 1-800-800-7441 or visit www.sgi.com/contactus.

Global Sales and Support: sgi.com/global