

University of Florida Case Study

The Interdisciplinary Center for Biotechnology Research at the University of Florida uses vSMP Foundation for SMP to provide high-performance, low cost SMP infrastructure



ScaleMP's vSMP Foundation platform for server aggregation solution provides low cost, dynamic virtualization solutions for flexible high performance workloads



BUSINESS BENEFITS

- Reduced capital expenditure outlays due to the avoidance of purchasing additional proprietary SMP systems.

IT IMPROVEMENTS

- A reliable, high-performance, scalable compute infrastructure.
- Faster application performance using virtual SMPs created with ScaleMP
- The ability to develop and iterate analysis parameters and take the time

Objective:

The Interdisciplinary Center for Biotechnology Research (ICBR) at the University of Florida (UF) needed a large shared memory system to continue to scale their legacy software, as well as the proprietary software packages that it used on a day-to-day basis. The goal was to find a way to expand its infrastructure to run compute intensive workloads as cost-effectively as the center's previous open source virtualization deployments.

Approach:

ICBR initially stretched its virtual infrastructure to capacity before investigating various single system image implementations available on the market. This evaluation led them to select ScaleMP's vSMP Foundation for SMP. This solution addressed the center's technical, as well as budget, requirements.

Customer Background:

The University of Florida (UF) is a public, land-grant, research university. The state's oldest, largest and most comprehensive university, UF has a long history of established programs in international education, research and service. The Interdisciplinary Center for Biotechnology Research (ICBR) at UF is a world-class research support center that provides biotech research services to the University of Florida faculty, staff, graduate students and other research partners throughout the state and nation.

A few years ago, ICBR created a computational infrastructure group that not only had an understanding of the biotechnology research being done

at the school, but also of the instrumentation and technologies required to support it. Aaron Gardner, an IT expert at the ICBR of UF, and four other IT team members support research at UF and abroad in various biotechnology fields such as proteomics, genomics, bioinformatics and cellomics. In addition to maintaining existing infrastructure, Gardner and his team strive to fill holes in their computational infrastructure to meet the needs of ICBR as well as other relevant researchers on campus. In order to complete projects they also rely on their partnerships with other data center organizations such as the high performance computing (HPC) center – where ICBR prefers to perform most of its compute intensive projects.

The UF campus has multiple campus data centers and a distributed IT backbone, but is currently formalizing a more centralized core structure. Gardner and his team at the ICBR data center maintain working relationships with all UF data centers including the HPC center and several other specialized labs and researchers. For local computation, ICBR's data center has around 100 servers for storage, virtualization, HPC and cluster computing capabilities. ICBR uses the Xen hypervisor in production and uses Red Hat Enterprise Linux across its systems, preferring open source options to proprietary solutions. In the future, ICBR plans to explore KVM and other emerging virtualization and cloud computing technologies.

CUSTOMER SOLUTION QUICK VIEW

Hardware Stack

- 16 x Intel Servers
- Processor: 32 x Xeon X5560 @ 2.8GHz (total 128 cores, 1.4 TFlops peak)
- Memory: 768 GB shared RAM

Virtualization Layer

- vSMP Foundation

Software Stack

- Red Hat Enterprise Linux (RHEL) 5
- Newbler
- MPI BLAST
- Home-grown applications

Customer Challenges

ICBR had many CPUs and plenty of computational power but could not harness this existing infrastructure to leverage its legacy software, as well as the proprietary software packages that it uses on a day-to-day basis. Gardner and his team needed access to a large shared memory system for performing these types of jobs that would not necessarily parallelize well.

ICBR needed a local system that would allow scientists and researchers to submit large interactive jobs. However, the price point of large SMP systems was prohibitive to ICBR. Because the large SMP systems the university looked at were all very expensive, Gardner and his team initially tried to find other ways to perform jobs. They ended up stretching their virtual infrastructure to accommodate these large shared memory workloads and using their infrastructure in ways that resulted in them beginning to lose the benefits of virtualization. The team placed larger and larger amounts of memory per node on HPC nodes, making it no longer cost-effective, especially given that non-HPC type jobs also performed on the same hardware.

Solution Selection Process

ICBR first learned about ScaleMP through Web research. Gardner had been investigating single system image implementations available on the market that would take commodity hardware and aggregate it into one operating environment with a large amount of shared memory. All of the existing technologies had limitations: many of them could aggregate, but would not allow multiple threads in a process to address the entire shared memory pool.

ICBR evaluated other approaches to pooling memory and investigated other vendors, but were unable to find a solution that does what ScaleMP does – create a true virtual SMP for a single system image in which a user can do anything on that system image that can be done on a traditional Linux-based SMP. ICBR also chose vSMP Foundation for SMP because it provides the high performance SMP solution the center was looking for and there was no additional development work that needed to happen in order to get the system up and running.

ICBR began deploying vSMP Foundation in mid-2009 in the machine room where all local computational resources are kept. The initial deployment took two days: one day to rack the hardware and set up, and the second day to install the vSMP Foundation for SMP software and Linux.

With vSMP Foundation, The University of Florida gains the following key benefits:

Although The Interdisciplinary Center for Biotechnology Research at the University of Florida is just beginning to move applications onto ScaleMP's vSMP Foundation-based systems, it has already realized or anticipates the following key benefits:

- Faster application performance using virtual SMPs created with ScaleMP.
- Decreased management complexity as opposed to traditional virtualization or cluster/aggregation technologies.

- The ability to approach next generation sequence processing in a way it could not do before; now ICBR can iterate on parameters and take the time to hone results versus having to accept whatever output it can get, which was previously the case because it took so long to get results.
- The option of moving more analysis software onto vSMP Foundation-based systems and using ScaleMP virtualized SMPs as an aggregation point for that type of work.
- The option of moving over test analysis workflows onto high-performing vSMP-Foundation-based systems to enable better long term success.