Voltaire
The Grid Backbone™

Case Study

Life Science Research Center Turns To Voltaire

For More Reliable and Faster Modeling and Analysis Throughput

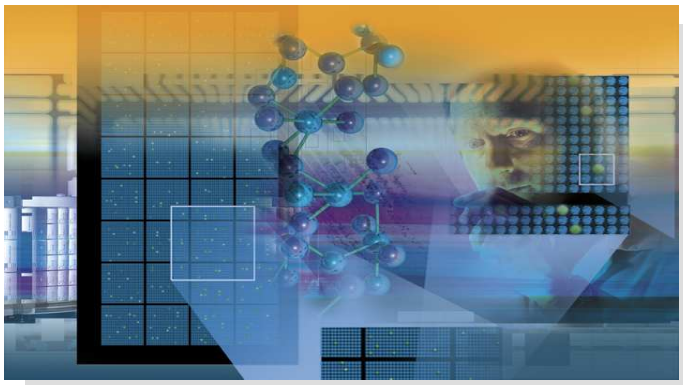

Swiss Institute of Bioinformatics Background

The Swiss Institute of Bioinformatics (SIB), an academic not-for-profit foundation, promotes research and the development of databanks and computer technologies in the bioinformatics field. SIB comprises researchers at universities and corporations across Switzerland and collaborates with researchers and institutions throughout Europe.

SIB also partners with industry leaders in the technology sector such as Hewlett-Packard, Intel and Oracle, which all helped SIB establish the high-performance Vital-IT Computer Center. The center is dedicated to life sciences, and at any one time, a dozen groups are conducting research projects at the center.

On behalf of researchers who utilize the computer center, Vital-IT focuses on two main missions: develop, optimize and parallelize life sciences software; and provide the necessary computing power for every research project.

“We are to an extent a software development shop since we optimize code,” said Dr. Victor Jongeneel of SIB. “This is useful to our industry partners who can show how well their software runs on the platform. At the same time, researchers come to us with large-scale analysis problems in all fields of life sciences because we can provide the horsepower and assistance they need for code development.”

The SIB Challenge

SIB first built the Vital-IT Computer Center in 2003 using a configuration relying on SAN storage and NFS services, using gigabit Ethernet for communication among servers. But 18 months after the initial deployment, SIB realized two major problems.

“With SAN storage, a limited number of servers can be connected to the SAN,” said Jongeneel. “Our servers send and receive data from the compute nodes using NFS, and this turned out to be a major performance bottleneck. With 64 clients running significant I/O, we needed a faster way than gigabit Ethernet for all of our compute nodes to share common file space.”

SIB also has many computing jobs that read large amounts of data into memory, many of which are I/O bound. Because of the load on the file servers, they would sometimes crash, and jobs could be aborted before completion. “Because of this, we could not run as many jobs per unit time as we wanted,” Jongeneel said.

These circumstances hampered researchers using Vital-IT that depended on the jobs to run to completion within a certain period of time. SIB also has a set of users that run parallelized programs using MPI libraries such as HP-MPI, Voltaire MPI and LAM MPI. These protocols allow message-passing between processes running on different nodes of a cluster.

Key Performance Results: For the SIB High-Performance Computing Center Powered by a Voltaire InfiniBand Switch

- Increased application and file-system performance: 10x better performance on HP SFS
- Lower latency
- Scalable to thousands of compute nodes
- Alleviated storage bottlenecks: fast file and block I/O
- Nearly 4 million jobs and more than 6 million atomic jobs run
- 784,628 CPU hours
- Increased MPI performance
- Almost 300 users
- 13 TB available disk space

Integrated Applications Running on the SIB Cluster Powered By Voltaire, Platform Computing and HP:

- Platform LSF
- AutoDock
- AutoFact
- BioConductor
- BIONJ - BLAT
- CHARMM
- Clustalw-mpi
- EMBOSS
- DOCK
- ESTScan
- Fasta/Fasty
- Garli
- GOLD
- HMMER
- IQPNNI
- LAMMPI
- Meme
- MERLIN
- Mpiblast
- mrbayes
- muscle
- NCBI BLAST
- Nemo
- PaCE
- Paml
- Pftools
- Phylip
- Phyml
- Probcons
- procse none
- QTDT
- RAxML V/VI
- T-Coffee TREE-PUZZLE

"There are a number of applications to parallelize using MPI," Jongeneel said. "Dependencies are created between programs running on different machines. They have to pass information back-and-forth to run large programs in a cluster environment. The performance of MPI on our system was poor because of the gigabit Ethernet connectivity and the magnitude of the applications running on the system." (see sidebar for list of applications).

Thus SIB began the search for server interconnects with higher bandwidth and lower latency.

The Voltaire Solution

SIB first considered a solution from a vendor featuring PCI-X throughput technology that met the technical specifications, but which also proved to be cost-prohibitive. SIB then turned to Hewlett-Packard for advice on how to increase bandwidth and lower the latency of the server interconnects. "There are not that many companies that offer a viable, affordable solution," Jongeneel said. "We decided to rely on the advice of HP as one of our trusted partners in the Vital-IT consortium."

Hewlett-Packard recommended SIB deploy InifiniBand-based Grid Backbone™ technology from Voltaire. "In addition to fitting within our budget, the Voltaire InfiniBand technology offered throughput that would more than double the speed of PCI-X technology," Jongeneel said.

Hewlett-Packard deployed two development clusters and two production clusters connected via a Voltaire Grid Director™ ISR 9288 switch. The system supports more than 13 terabytes of disk storage and provides more than two Tflops of throughput at peak performance. Both production clusters contain 32 nodes, with one utilizing Itanium2 dual processors. The other cluster utilizes Xeon dual processors, delivered as HP Cluster Platforms, and running Red Hat Enterprise Linux.

"Voltaire is our preferred supplier for InfiniBand technology in our HP Cluster Platforms," said Ed Turkel, Manager Product and Technology Marketing for Hewlett-Packard. "Through our partnership with Voltaire, we are able to provide a wide range of solutions across our extensive product family, providing industry-leading performance for our customers' most challenging science, engineering and analysis problems."

Jongeneel added, "For this configuration, the technology most clearly supported by HP was InfiniBand. It surpassed our technical specifications for bandwidth and latency."

The Results

"The Voltaire interconnect between compute nodes provides much better performance both in terms of bandwidth and latency," Jongeneel said. "Computing jobs using HP-MPI, Voltaire MPI and LAM MPI now run much faster than they used to, and more importantly, our server I/O capacity is no longer a bottleneck for running any size job."

The Voltaire technology also allowed SIB to implement a cluster file system providing direct data access to the compute nodes, a critical feature for SIB because this eliminates performance, availability and scalability problems present in many traditional, distributed-file systems. SIB chose to implement HP StorageWorks Scalalable File Share (SFS) to solve their data access issues.

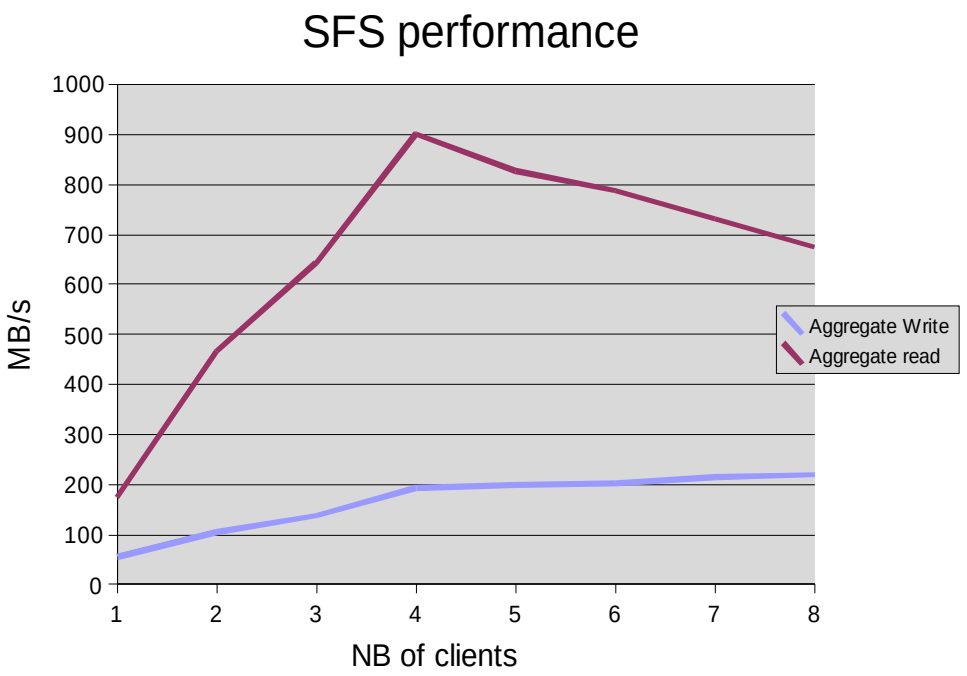

SIB has found that the combination of the Voltaire Grid Director™ ISR 9288 switch and HP SFS gives researchers I/O performance comparable to locally-attached discs. Jongeneel described one particular project where data is passed back-and-forth through the file system between a genetic algorithm that assigns metrics and positions to a set of atoms, and another program that performs molecular minimization and molecular dynamics on the same set of atoms.

"There's a constant back-and-forth exchange of data between the two programs that access the same data," Jongeneel said. "We have seen significant throughput improvement since installing the Voltaire switch in January of 2006. Overall, the HP and Voltaire cluster has increased our file system throughput by 10 times, and our researchers have doubled their productivity in running analyses and simulations."

VOLTAIRE
The Grid Backbone™

Contact Voltaire

1.800.865.8247 info@voltaire.com www.voltaire.com