

Solution Brief:

Higher Education and Research

Bringing New Levels of Performance to High Performance Clusters



Contents

- The Higher Education and Research Lab Challenge** 1
- Today's Solution** 2
- A Better Way** 3
- Building High Performance Clusters** 4
 - High-performance InfiniBand Switches** 4
 - Fast Storage Connectivity** 6
- Putting It All Together** 7
 - Key Features & Benefits** 7
 - Tested & Certified with Leading Applications** 8
 - Industry Standard Performance Benchmarks** 9
 - LINPACK** 9
 - SPECenvM2002** 10
- Voltaire Higher Education and Research Lab Customers** 11
- About Voltaire** 12

Abstract

Higher educational institutions and research labs constantly strive to uncover and apply important new discoveries. The intellectual talent and computing capabilities of these labs are used by government agencies and corporations to help solve complex problems or to accelerate research and development projects. This requires a computing environment that can run complex applications as fast as possible with the best possible price to performance ratio. To dramatically improve the speed and efficiency of their server-based clusters and storage systems, higher educational institutions and research labs are turning to Voltaire's InfiniBand-based solutions.

The Higher Education and Research Lab Challenge

The need to run applications faster: Many educational institutions are contracted by government agencies, manufacturers, financial services companies, and media corporations to partner in research. They help corporations develop their products faster by providing access to their expertise and their high performance systems. They also help government agencies and corporations speed up their product life cycles to help them bring products to market faster with more accuracy to pinpoint problems and achieve better quality.

Race for discovery in their research area: Higher education and research labs are tackling complex problems that require significant computing capabilities. Having systems that solve problems faster enables them to be the first to discover new and exciting breakthroughs in their areas of expertise.

The desire to sell cluster time to outside companies: Higher educational institutions work closely with government agencies and corporations to perform product discovery and research. If the institution or lab can aid a corporation in developing a product faster and with better quality, the sub-contracting jobs will keep returning.

Increase productivity: These institutions and labs need to make the most of their limited resources. This requires an increase in the productivity of the clusters they are running. If they are able to increase the performance of their cluster even 10%, they can be that much ahead of their competition in making discoveries within their areas of expertise.

Solution Brief: Higher Education and Research

Today's Solution

Large symmetric multi-processing machines (SMPs) used to be the answer for generating compute power in the data center. However, these proprietary, expensive systems gave way to cluster and grid architectures consisting of low-cost commodity elements that offer comparable performance.

Some data centers try to scale clusters by adding additional servers or moving to servers with multiple cores. This approach can work for smaller, simpler simulations, but the more complex the analysis, the more likely the need to run simulations across multiple servers where latency is a major factor in determining performance. The answer to speeding analysis and maximizing return on supercomputing investments is not simply buying more or bigger servers, but rather eliminating bottlenecks to performance by employing the use of a high performance interconnect.

Because of the ready-availability of Ethernet, many of today's clusters are built with Ethernet as the interconnect. While Gigabit Ethernet-based clustering is cheaper than SMP-based architectures, it can be very inefficient. For applications that rely on bandwidth or memory sharing, the efficiency (percentage of a server-CPU dedicated to communications overhead) can be a concern.

Today, many high performance data centers take advantage of InfiniBand-based interconnect support incorporated into leading ISV applications and MPI-based applications. Voltaire InfiniBand interconnects eliminate I/O bottlenecks, allowing applications to run faster and more efficiently.

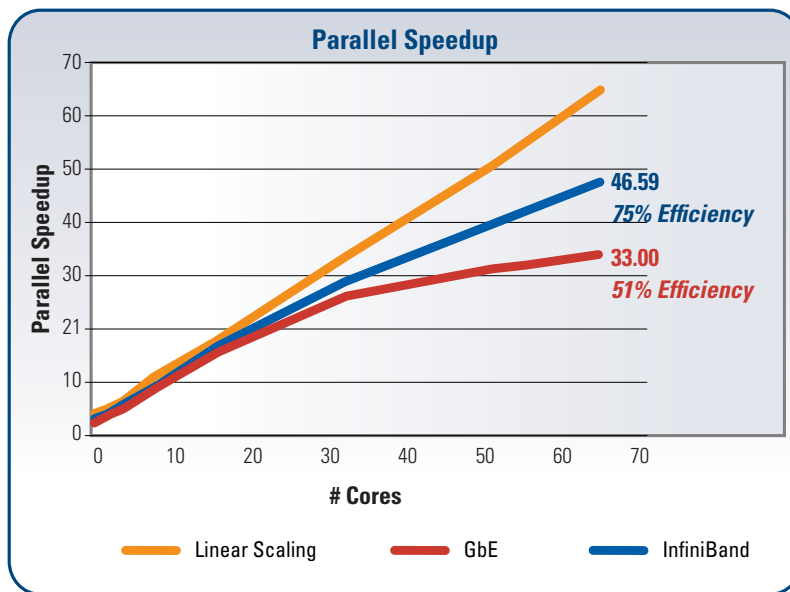


Figure 1. Voltaire InfiniBand-based solution improves performance by 50%

A Better Way

Voltaire InfiniBand solutions make applications run faster. Voltaire offers high-performance (10, 20 and 40 Gbps), low-latency (< 2 microseconds) interconnect solutions for high performance data centers. Benchmark testing has found that Voltaire interconnect solutions reduce application runtime by as much as 50 to 300 percent.

In addition to offering InfiniBand switch-technology, Voltaire works directly with software-vendors to create the most efficient, fastest, and lowest-latency solutions available. By leveraging Voltaire InfiniBand-based solutions, organizations can now speed application performance to analyze products faster and more efficiently.

As today's price and performance leader in the industry, Voltaire builds its solutions using standards-based InfiniBand technology. InfiniBand is an industry-standard interconnect for high-performance computing (HPC) and enterprise applications. The combination of high bandwidth, low latency and scalability makes InfiniBand the interconnect of choice to power many of the world's largest and fastest computer systems and commercial data centers. Voltaire solutions support most major server vendors, operating systems, storage solutions and chip manufacturers.

Benchmark testing has found that Voltaire interconnect solutions reduce application runtime by as much as 50 to 300 percent.

	1 Gb Ethernet	10 Gb Ethernet	Myrinet	InfiniBand
Bandwidth	1 Gb/sec	10 Gb/sec	2.5 Gb/sec	10, 20 & 40 Gb/sec
Latency		~10 us	2.5 - 5.5 us	< 2 us
Average Efficiency	53%	No Entries	68%	74%
Price Per Gig/Port	~\$350.00	>~\$700.00	~\$225.00	<\$100.00

Table 1: Price/performance advantages for InfiniBand

Building High Performance Clusters

Voltaire offers complete end-to-end server interconnect solutions for speeding applications. The two major elements of the solution include:

- High-speed, low latency InfiniBand switches
- Fast storage connectivity

High-performance InfiniBand Switches

Voltaire's InfiniBand-based solutions deliver high performance and scalability to compute clusters. Voltaire offers a complete portfolio of products including a scalable line of InfiniBand switches, high performance I/O gateways (for seamless connectivity to Ethernet and Fibre Channel networks) and fabric management software. Voltaire solutions use the Open Fabric Alliance's OFED drivers and the Open MPI (Message Passing Interface) libraries to optimize application performance for both MPI-based and socket-based applications.




Figure 2. Voltaire Grid Director 9024 for small-to-medium sized clusters ranging from 16 to 24 nodes

For small-to-medium sized clusters, Voltaire offers the Voltaire Grid Director™ 9024. It is a 1U device with twenty-four 10 Gbps (SDR) or 20 Gbps (DDR) InfiniBand ports. The switch is a high performance, low latency, fully non-blocking edge or leaf-switch with a throughput of 480 Gbps.

It is well-suited for small InfiniBand fabrics with up to 24 nodes because it includes all of the necessary management capabilities to function as a stand-alone switch. The Grid Director 9024 is internally managed and offers comprehensive device and fabric management capabilities. Designed for high-availability (high MBTF) and easy maintenance, the switch is simple to install and features straightforward initialization. The solution is scalable as additional switches can be added to support additional nodes.




Figure 3. Voltaire Grid Director 2004 for scalable clusters ranging from 24-96 compute nodes.

For larger clusters ranging from 25-96 compute nodes, Voltaire offers the Grid Director™ 2004 multi-service switch — the industry's highest performing multi-service switch for medium-to-large clusters and grids. The switch enables high performance non-blocking configurations and features an enterprise-level, high availability design. The Grid Director 2004 supports up to 96 InfiniBand 4X ports (20 Gbps) and is scalable through the use of additional, hot-swappable modules. The Grid Director 2004 also features 10 GbE and Fibre Channel ports so the solution can provide high-performance, integrated SAN and LAN connectivity.

For very large clusters and grids, Voltaire offers the Grid Director™ 2012, a larger version of the Grid Director 2004. The Grid Director 2012 features integrated storage and LAN connectivity, and it supports up to 288 InfiniBand 4X ports (20 Gbps).

Solution Brief: Higher Education and Research

Fast Storage Connectivity

For companies looking to incorporate storage into their InfiniBand clusters, Voltaire solutions offer Fast I/O capabilities for storage. Voltaire solutions combine scalable compute and storage capabilities by using parallel file systems. Running scalable file systems over Voltaire InfiniBand solutions creates the most scalable solution in the industry-with more than 1000 nodes on a single name space-and delivers high performance connectivity for the storage and client nodes.

Manufacturing applications can now have effective File I/O rates of 350MB/s compared with the 50MB/s previously available using NFS.

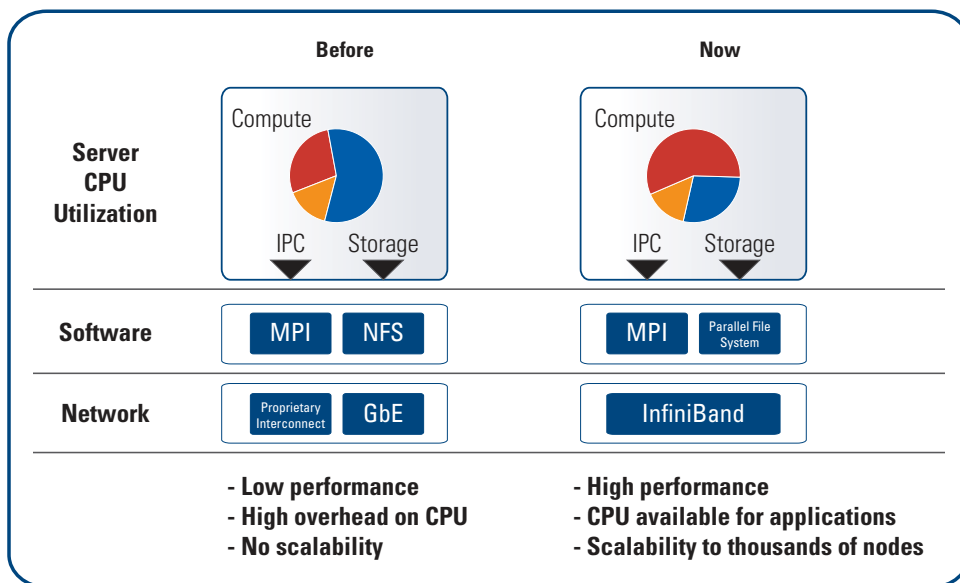


Figure 4. The advantages of combining IPC and File I/O with Voltaire Solutions.

At the heart of the solution is the Voltaire Grid Director 2004 multi-service switch (described in the above "Scalable Solutions" section). Voltaire's director-class, multi-service switches offer seamless InfiniBand, GbE and Fibre Channel connectivity. This enables MPI and storage traffic to run on the same network, a capability Ethernet and proprietary fabrics do not offer. By enabling IPC and high-performance storage on a single network, Voltaire solutions allow companies to leave behind the limitations of network file systems (NFS) and move to parallel file systems over InfiniBand. This provides far greater scalability.

Manufacturing applications can now have effective File I/O rates of 350MB/s compared with the 50MB/s previously available by using NFS. Additionally, the size of compute clusters is no longer limited because of the limitations imposed by NFS.

Putting It All Together

Key Features & Benefits

Voltaire solutions for higher education and research offer many compelling benefits to users:

- **High Bandwidth For Faster Product Development and Problem Solving:** Voltaire solutions provide bandwidth of 20 Gbps to allow for faster and more frequent analyses by commercially available or home-grown applications.
- **Lower Latency To Accelerate Application Runtime:** Voltaire solutions provide latency as low as 1.3 microseconds. This reduced latency over Ethernet technologies will accelerate the runtime of applications for labs that are sharing resources, or trying to help an outside agency or corporation solve a specific problem.
- **Standards-Based and Flexible:** Voltaire solutions are based on InfiniBand: the only industry-standard, high-performance interconnect for grids and clusters. Voltaire solutions can be built as fully non-blocking 20Gbps, or as a lower-bandwidth fabric based on the needs of the applications. This is important to anyone developing in an environment that uses multiple commercial or homegrown applications. The easier and more flexible the applications, servers, and interconnect environment, the faster the organization will be productive solving problems, sharing data, and making important discoveries.
- **Fast I/O for Storage:** Voltaire solutions enable parallel file systems over InfiniBand, which offers far-greater performance and scalability than NFS solutions.

Solution Brief: Higher Education and Research

Tested & Certified with Leading Applications

By working closely with leading server and software vendors on integration and testing, Voltaire offers the fastest and most efficient high-speed interconnect solutions for most leading applications.

Voltaire solutions support the following applications, MPI offerings, management applications, operating systems and parallel file systems.

Applications	Abaqus-Abaqus, Ansys-Ansys, CD-Adapco-StarCD, ESI-Group-PamCrash, ESI-Group PamFlow, EXA-PowerFlow, Fluent-FLUENT, LSTC-LS-Dyna, Mecallog-Radioss, MSC Software-Nastran, Ricardo-Vectis
MPI Support	OpenMPI, MVAPICH, HP-MPI, Intel MPI, SGI MPT, Scali MPI Connect, ParaStation MPI, Microsoft MPI (MS CCS)
OS Support	Linux Enterprise Edition from Novell (SUSE SLES) and Red Hat (EL/AS)
Supported Parallel File Systems	Lustre, HP SFS, IBM GPFS, Panasas, IBRIX, TerraScale's TeraGrid
Systems & Platform Partners	HP, IBM, SUN, NEC, SGI, NEC, Intel, AMD



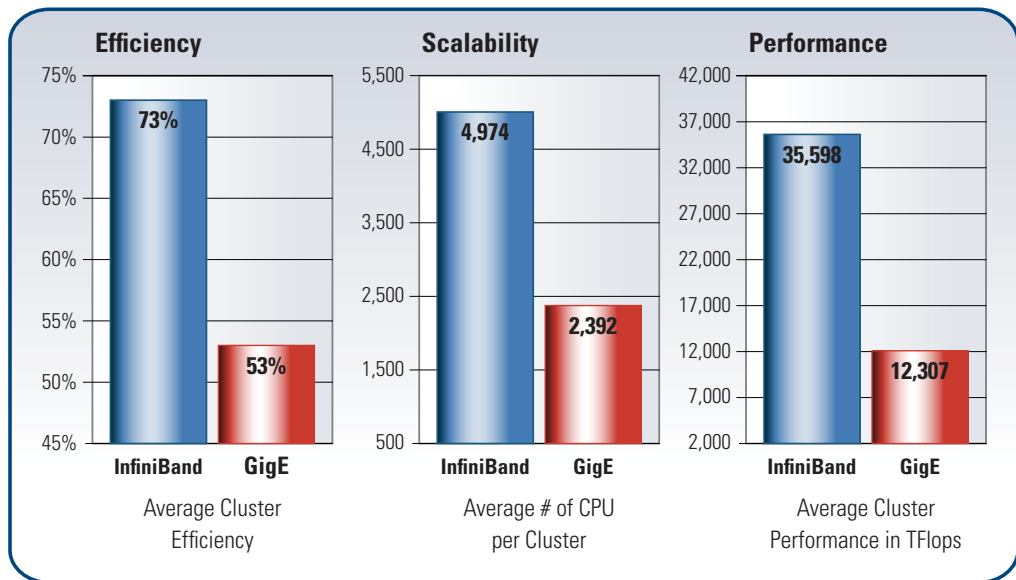
Industry Standard Performance Benchmarks

Application level performance benchmarks reinforce the advantage of InfiniBand. InfiniBand leads by a substantial margin among many industry standard benchmarks relevant to a wide variety of industries and disciplines.

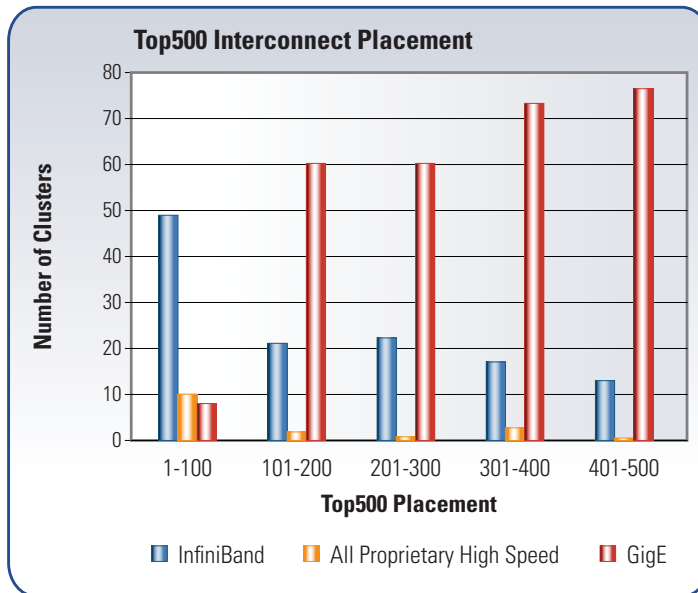
LINPACK

The Linpack benchmark measures floating point performance and memory bandwidth by solving a dense system of linear equations. It allows the user to scale the problem size and optimize the software in order to achieve the best performance for a given machine.

Linpack performance does not reflect the overall performance of a given system, as no single number ever can. It does, however, reflect the performance of a dedicated system for solving a dense system of linear equations. Since the problem is very regular, the performance achieved is quite high, and the performance numbers give a good indication of peak performance. This benchmark is used in the ranking of Top500 clusters.



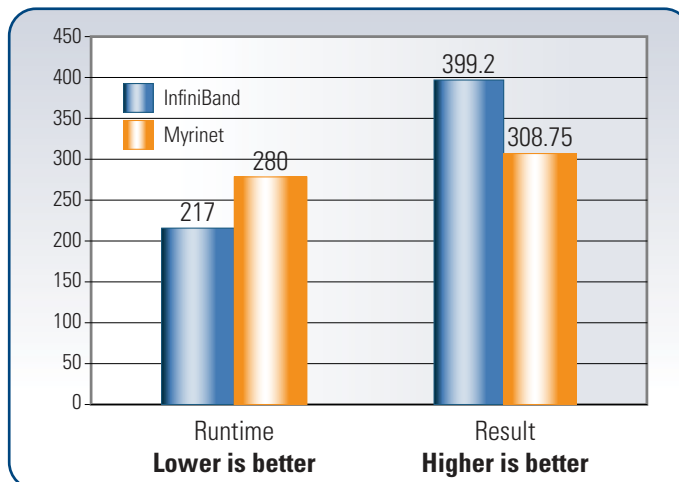
Solution Brief: Higher Education and Research



SPECenvM2002

The SPEC ENV2002 benchmark is based on the Weather Research and Forecasting (WRF) model, a state-of-the-art, non-hydrostatic mesoscale weather model. SPEC/HPG integrated version 1.2.1 of the WRF model into the SPEC tools for building, running and verifying results.

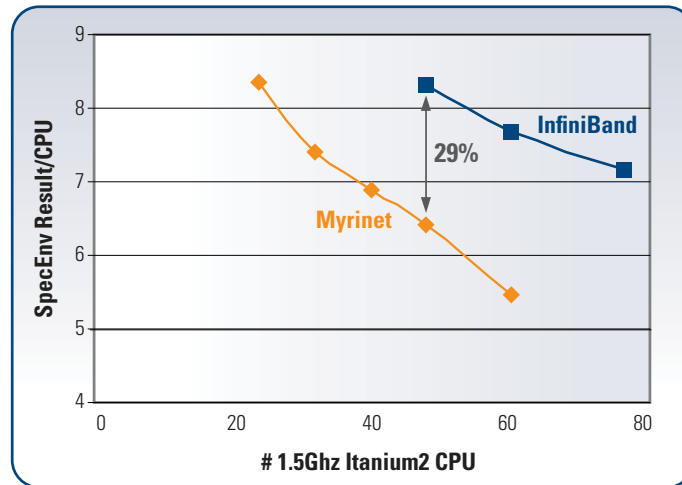
The SPECenvM2002 chart below shows the performance of a computing system in simulating the weather over the continental United States for a 24 hour period at a 22km resolution using the WRF Model.



On Identical Platforms the SPECenv results reveal InfiniBand performance over Myrinet


This performance benchmark of real MPI-based HPC applications runs on identical platforms and shows performance with both Myrinet and InfiniBand interconnects. It provides precise apples-to-apples interconnect comparison and demonstrates the performance benefits of InfiniBand.

The chart below from the same benchmark shows the scalability advantage with a 29% improvement for 48 CPUs; the gap increases even more to 40% for 60 CPUs. As the cluster grows and more CPUs are added to the system, the gap continues to widen.



The SpecEnv performance per CPU reveals increasing advantage and efficiency as the cluster CPU count increases.

Voltaire Higher Education and Research Lab Customers



About Voltaire

Voltaire (NASDAQ: VOLT) designs and develops server and storage switching and software solutions that enable high-performance grid computing within the data center. Voltaire refers to its server and storage switching and software solutions as the Voltaire Grid Backbone™. Voltaire's products leverage InfiniBand technology and include director-class switches, multi-service switches, fixed-port configuration switches, Ethernet and Fibre Channel routers and standards-based driver and management software. Voltaire's solutions have been sold to a wide range of end customers including governmental, research and educational organizations, as well as market-leading enterprises in the manufacturing, oil and gas, entertainment, life sciences and financial services industries. More information about Voltaire is available at www.voltaire.com or by calling 1-800-865-8247.

Notice

Reproduction of this publication in any form without prior written permission is not allowed. The information in this publication is subject to change without notice and is provided "AS IS" WITHOUT WARRANTY OF ANY KIND. THE ENTIRE RISK ARISING OUT OF THE USE OR INTERPRETATIONS OF THIS INFORMATION REMAINS WITH RECIPIENT. IN NO EVENT SHALL VOLTAIRE BE LIABLE FOR ANY DIRECT, SPECIAL, PUNITIVE OR OTHER DAMAGES.

Performance results will vary based upon a number of system factors. Some of these include: server configuration of the processor, chip set, memory size, firmware and driver release versions, MPI version and OS kernel version. The configuration or configurations tested or described may or may not be the only available solution. These tests are not a determination of product quality or correctness, nor does it ensure compliance with any federal state or local requirements.

Product names mentioned herein may be trademarks and/or registered trademarks of their respective companies.



Contact Voltaire to Learn More

1.800.865.8247
info@voltaire.com
www.voltaire.com

©2008 Voltaire Inc. All rights reserved. Voltaire and the Voltaire logo are registered trademarks of Voltaire Inc. Grid Director is a trademark of Voltaire Inc. Other company, product, or service names are the property of their respective owners.