



Accelerating Simulations at Tech-X with the Terascale High Throughput Storage Appliance




Terascale, Inc.
145 Bodwell St
Avon, MA 02322
www.terascal.com

Overview

Tech-X Corporation of Boulder, Colorado develops software solutions for research, engineering and education to aid with specific scientific challenges. Tech-X Corporation offers commercial products as well as collaborates with government and private institutions on research projects. Products and services from Tech-X help solve difficult technical problems, increasing design and development productivity and accelerating project deployment for scientists and engineers.

Summary

Tech-X recently selected the Terascale RTS 1000 storage appliance to use in code development for its VORPAL software and to run simulations using VORPAL for Tech-X customers around the world. VORPAL is a cross-platform framework capable of large-scale, massively-parallel simulations of electromagnetic and electrostatic phenomena, modeling fields, particles, and fluids in one, two, or three dimensions. VORPAL offers a unique combination of physical models to cover the entire range of plasma simulation problems. Plasma processing chambers, plasma thrusters, particle accelerators, and high-power microwave guides are only a few of the many applications benefiting from the powerful, parallel algorithms incorporated into the VORPAL framework. World-class research organizations such as Tokyo Electron America, Inc, Lockheed Martin, Lawrence Berkeley National Laboratory, Brookhaven National Laboratory, and the University of Colorado at Boulder are actively using VORPAL in research. VORPAL is designed to run as both a serial code for single-processor workstations and as a parallel code for systems that support MPI. Domain decomposition techniques employed in the code enable static and dynamic load balancing. Strong scaling on several thousands of processors has been achieved for VORPAL simulations. The use of parallel I/O yields high performance even when thousands of processors are dumping large amounts of data.



Searching for a Higher Performance Solution Capable of Running Large Simulations

Anne Hammond, Director of Information Technology at Tech-X, is responsible for Tech-X's IT strategy and for building a core set of systems capable of running large simulations using the company's VORPAL software. Tech-X had been running simulations using an NFS system but wanted to find a solution that would enable the company to run larger and longer simulations for both internal development and for corporate customers and scientific institutions worldwide.

"VORPAL has very flexible code and we wanted to improve our ability to run large simulations for customers. We really didn't have the ability to do the kind of simulations we wanted because we had a smaller cluster that wasn't capable of writing out very large data files," said Hammond. "We really needed to move up to a system that would enable us to perform larger and longer simulations."

Hammond and her team began evaluating parallel file systems and eventually chose the Terascale RTS 1000 after looking at hardware-based NFS, object-storage based, and SAN based competitive solutions. The company also considered using building its own solution but eventually deemed the process to be too complex.

"We looked at several different solutions but in the end, we found that the Terascale RTS 1000 was the ideal product for our environment. Some of the other systems we looked at were too large for our needs but the Terascale system was perfect in terms of throughput


and manageability. It also came in at a very attractive price point,” said Hammond. “It met our budget requirements, fit easily into our existing environment, and its appliance packaging helped us reduce the time it took to get the system up and running.”

Scalability was another important factor in choosing the Terascale RTS 1000. The system was designed to be easily scalable, and throughput can be easily increased by adding additional object servers.

“We knew we had a certain amount of funding for the project and we were able to tailor the system to meet our budget. It’s nice to have the option to add things like redundant metadata storage and redundant power supplies, and we will consider increasing functionality and throughput as our budget and our needs grow,” she said.

High Performance, Appliance Packaging for Easy Deployment and Management

“Some of the data files we generate can be huge and we need to perform multiple steps to complete the simulations. The Terascale RTS 1000 gives us the throughput we need to perform larger and longer simulations. The speed more than meets our expectations, with 32 simultaneous client node writes, using InfiniScale IV QDR InfiniBand interconnects, we are experiencing ~650 MB/sec without any tuning of the application.” Hammond said.



Simplicity was a key reason Tech-X chose the Terascale RTS 1000. It is delivered as an appliance and was designed to be simple to deploy and manage. The RTS 1000 comes with all software installed and is tuned for high performance.

“We liked the fact that the RTS 1000 came delivered as a working appliance. We have a small, agile staff and we wanted the system to be fully operational as soon as possible. We were able to get the system up and running very quickly and didn’t have to worry about the technical details. It just worked,” said Hammond. “There’s a myth out there that a parallel file system

is far more complicated to operate and manage than a standard NFS box. That certainly hasn’t been the case with the Terascale system. In fact, we’ve found the Terascale RTS 1000 to be extremely easy to manage. It works smoothly and our staff doesn’t have to invest a lot of time in the system. It enables us to focus on our applications and so that we can better support our users,” she said.

Lustre for Ease of Use

Hammond noted that from a user standpoint, the fact that the Terascale RTS 1000 is based on Lustre makes it easy for users to get up to speed quickly.

“Our users were able to get a jump start on getting their application to run faster because we didn’t have to concentrate on building the Lustre infrastructure for them. It really helped make our users more efficient because they were able to go on and take advantage of the system quickly without waiting for us,” she said. “For example, there’s only one command to make a directory that allows a user to write in parallel, so our users could easily test their application under their own directory. They were able to come up to speed with very little user education. All in all, the Terascale RTS 100 helped make our users more efficient and our IT staff more effective.”

Terascala's Storage Appliance

The Terascale RTS 1000 Storage Appliance is a Lustre parallel file system-based system that offers high throughput and high capacity. The RTS 1000 is designed to deliver the maximum throughput to enable applications to run at peak efficiency. It can deliver over 2GB/sec from a single enclosure and up to 10 GB/sec for a full rack solution. Designed to plug directly into the compute client network environment, the RTS 1000 has an optimized data path from the client network through to the disks within the storage device.

Simplification of deployment and ongoing management is a key aspect of the RTS 1000. It is delivered as an appliance with all of the software installed and tuned to deliver performance. With its built-in management system, it is easy to add additional capacity and throughput or to fine tune key aspects of the environment. The RTS 1000 is designed to be operated and managed by system administrators without extensive storage experience.

Terascale leverages Lustre because it is an open source, high performance clustered software initially developed for applications needing very high throughput, scalability and capacity. It offers high levels of reliability, scalability and performance, having been deployed



in some of the largest compute installations in the world. Lustre leverages a simple metadata/stored object architecture, where the metadata server stores location information about data and the object store servers act as the repository for the actual data. This approach allows throughput to be scaled by simply adding additional object store servers. Terascale has an optimized architecture and has tuned Lustre for the specific capabilities of the RTS 1000, so the appliance delivers a simple, easy to use and expandable solution.

With its simplified architecture and the use of commodity technologies where possible, the RTS 1000 is able to deliver tremendous price/performance while simplifying the overall deployment and management of a high throughput storage solution.

The RTS 1000 is part of a family of parallel file system based appliances from Terascale. All the Terascale solutions include the Terascale Integrated Storage Information System (TISIS™) which provides real time storage optimization and complete management. Terascale solutions are all delivered as fully configured, ready to go appliances. Contact Terascale at www.terascala.com or info@terascala.com to learn more about all our solutions.



Terascale, Inc.
145 Bodwell St.
Avon, MA 02322
Tel: 508-588-1501
Email: sales@terascala.com
www.terascala.com

© 2011 Terascale, Inc. Terascale is a trademark of Terascale, Inc. Lustre is a registered trademark of Oracle. All other trademarks are the property of their respective holders.