



High Throughput Storage at a Department of Energy Laboratory



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

Summary

A major government laboratory focused on developing nuclear weapons, energy and infrastructure assurance, non-proliferation, and defense systems and assessments recently selected the Terascala™ RTS 1000 Storage Appliance to provide high throughput storage for new research initiatives.


Background

The scientific computing team at the laboratory was in the process of building a new Linux cluster and needed a high performance, high throughput storage solution. The team had worked with and evaluated Lustre™ in the past but became frustrated with the difficult implementation process and inconsistent performance results. In the end, the team deemed Lustre “not ready for prime time.”

The laboratory had been using network attached storage (NAS) systems from a variety of different vendors but needed a solution that would provide higher performance and speed and decided to reconsider Lustre. But this time, the laboratory decided to evaluate the Lustre-based Terascala RTS 1000 storage appliance.

“We brought in the Terascala unit for evaluation and worked with it for several months. We were extremely happy with its performance. None of the other solutions we looked at came close in terms of performance,” said a scientific computing director at the laboratory.

“Price was also a big factor. The Terascala RTS 1000 provided us with the throughput and high performance we needed at a great price point.”



Terascala Meets High Performance Storage Challenges

The new system will be used by scientists at the laboratory who primarily work with finite element and matrix multiplication to solve a range of problems. For example, one application will demand that a single job be run on over 200 compute nodes, potentially generating in excess of 40 terabytes of data. Other jobs will generate hundreds of gigabytes to a terabyte of data that will be stored on the Terascala RTS 1000, with jobs run repeatedly to verify and validate results. The laboratory's RTS 1000 installation provides 120 terabytes of storage and delivers up to 6 GB/sec of throughput.

“We had been using NAS-based storage systems to support some of our runtime storage needs but it didn't provide the high performance and high speed we need for our most intensive applications,” said the director. “We believed Lustre was the answer but we didn't want to go through the pain of building our own system. The Terascala system is an appliance and is fully configured and tuned for high performance, high throughput storage and was exactly what we needed.”

Flexibility to Meet New Challenges

For the laboratory, scalability was also a key consideration in choosing the RTS 1000.

“Scalability is important for any technology purchase we make. We are at the cutting edge in terms of using technology to solve far-reaching, critically important problems. We need the flexibility to grow our systems and storage easily and seamlessly to meet our emerging needs, and Terascale’s architecture enables us to increase performance as our clusters grow,” said the director. “We purchased the Terascale system to use with an InfiniBand cluster but it’s nice to know that the RTS 1000 is flexible enough to use with Gigabit Ethernet or 10 Gigabit Ethernet.”

Terascale’s RTS 1000 architecture provides the ability to expand throughput or capacity by adding additional nodes without impacting the existing environment. Additionally, the RTS 1000 incorporates “plug and play” replaceable units, hot swappable components for reliability, redundant power supplies and high density packaging along with support for GbE, 10GbE or InfiniBand.

Terascale’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.

Contact Terascale at www.terascalacom or info@terascalacom to learn more.

About Terascale

Terascale develops software solution for high throughput, high capacity and cost-effective storage appliance. Its unique storage appliance approach is changing the dynamics of the performance driven computing market, enabling existing users to do more for less while enabling new users to maximize the capabilities of their processing infrastructure. Founded in 2005, Terascale is based in Avon, MA. Learn more at www.terascale.com



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