



TERASCALE

## High Throughput Storage at Arizona State University



Terascale, Inc.  
145 Bodwell St  
Avon, MA 02322  
[www.terascale.com](http://www.terascale.com)

## Summary

The Fulton High Performance Computing Initiative offers world class computing resources to the researchers and students in the Ira A. Fulton School of Engineering at Arizona State University (ASU). The Initiative is a hub for research in computing, collaborative research in the application of parallel computing systems, and a center for education in high performance computing systems. With over 65,000 students and 8,000 faculty members, the Initiative supports applications across 120 different faculty groups in 35 departments at the University.

The team at the Fulton Initiative recently installed the Terascale RTS 1000 storage appliance to provide high throughput storage for their research scientists.

## Background

One of the key roles of the Initiative is to assist with any research project that is intensive in terms of compute cycles or data set size. Working with such large group of applications and a varied user base is challenging, especially considering the quantity of data being processed and the amount of data growth that the Initiative experiences each year.

Dr. Dan Stanzione founded the Fulton Center and is currently its director. Stanzione and his team built the center from the ground-up, first building large computing clusters and then branching into storage, visualization and other services. Together, all of the Fulton Center's services combine to make a cyber infrastructure that enables researchers to work as productively as possible.

ASU's research projects are far-reaching and often require large scale, high speed file systems for researchers to be able to perform their work effectively. Three examples of this research are weather prediction, astronomy and astrophysics, and atomization.

ASU researchers working on weather prediction leverage environmental fluid dynamics and atmospheric physics to produce fundamental predictions on the structure of the atmosphere. The research generates very large amounts of weather simulation data.

"You have to do [weather simulation] quickly because there's nothing more useless than yesterday's weather forecast," said Stanzione. "For example, a 48 hour forecast in a 500 kilometer grid around Phoenix will consume almost a terabyte of data and you need to be able to move that data to storage quickly."


High throughput storage is critical in research on traditional computational fluid dynamics, specifically in the study of atomization. "We're looking at the dynamics of where a spray fluid turns into a bunch of little particles," said Stanzione. "That affects everything from how you design and formulate spray paint to how much carbon dioxide gets absorbed by the ocean in the surf to the efficiency of combustion in engines as you're spraying fuel into an engine cylinder." Atomization creates and demands a massive amount of data because the researchers are looking at the interaction of the spray and the surrounding environment, representing a huge number of particles. This type of research creates terabytes of data on a thousand processor simulation.

In the areas of astronomy and astrophysics, ASU researchers are gaining a better understanding of high energy physics by looking at colliding supernovas, work that can



*Dr. Dan Stanzione, Director,  
Fulton High Performance  
Computing Initiative, Arizona  
State University*

generate terabytes of data in the span of hours. The challenge is that the system must write the data quickly or the computation will slow down waiting for the storage system. This makes storage throughput essential. Once the data has been gathered, ASU researchers perform analysis on it quickly to produce movies / visualizations or datamine it to better understand the dynamics of the event. The real “information” needed is not in the raw data, but in the mined version, so the huge amount of data generated is only needed for a short time. These simulations will generate six to eight terabytes of data from a single computation that takes one to three days to run.



*Visualization of hierarchical volumetric data from simulation of black-hole driven bubbles. (Adam Kubach, Fulton HPC Visualization Team, ASU)*

### **Storage Challenges at ASU**

ASU's Fulton High Performance Center leverages hundreds of nodes of “commodity component” based servers that are linked together via a high speed network that enables them to work together to solve complex problems. While this approach has become the workhorse of today's supercomputing workloads, as compute nodes become more powerful it has been an industry-wide challenge to deliver storage solutions that deliver data quickly enough to maximize use of the environment.

“From the start, we sought out storage solutions that could keep pace with the growing amounts of compute power that we were building,” said Stanzione. “Storage throughput is extremely critical with large, supercomputing-type systems running scientific simulations. The system needs to have enough throughput speed to keep up with the processor.

“Supercomputers are expensive, and even with their million dollar price tags, they have a short lifespan and depreciate thousands of dollars each month,” he continued. “The last thing you want to do is take that system and have it spending half its time waiting for data to write out to the disk so you can keep the computation going. What is needed is a file system that can receive data as fast as the computer can generate it. That really is a key for us. We can't waste time on our million dollar supercomputer, and we can't waste the time of our researchers.”

### **How Terascale helps the ASU Fulton High Performance Initiative**

Stanzione and his team at ASU have been working with a number of storage solutions in the center and evaluated Terascale's RTS 1000 run time storage appliance when the University decided to look at solutions to complement its existing storage offerings.

“Our existing solutions run proprietary file systems and can be extremely expensive to purchase and to run. We looked at the Terascale solution in an effort to maintain high performance and minimize costs,” said Stanzione. “The fact that the RTS 1000 runs Lustre™ is an advantage because it puts us in line with what our users have seen at other centers. It helps us create a more consistent environment for the users across all different

systems, such as those at NASA, the Department of Energy, the Department of Defense, the National Science Foundation and through the TeraGrid project. While there's no such thing as a standard file system in supercomputing, Lustre is as close as there is to a de facto standard. From a programmability and application perspective, it's what our users expect."

### **Performance, Scalability, Cost-effectiveness and Support for Open Standards are Key**

Some of the key considerations Stanzione and his team looked at were overall performance, scalability and cost effectiveness, as well as support for open standards. ASU chose the Terascale solution after an extensive evaluation process and currently offers the solution to its researchers seeking high throughput storage.

"We performed fairly extensive benchmarking on the RTS 1000 and were very pleased with its performance," said Stanzione. "The RTS 1000 give us the performance our researchers demand to keep up with their storage needs and enables us to offer it cost-effectively. Storage is a major component of our IT spending and we are always trying to reduce costs in that area. The Terascale solution offers great price-performance."

Scalability is also important to ASU as the university's computing needs grow.

"Terascale's blade architecture allows us to very quickly scale out performance as our clusters grow," said Stanzione. "It's all about performance and throughput, and smooth scalability is critical for future applications."

The RTS 1000's blade 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 the choice of GbE, 10GbE or InfiniBand.



"Because it is an appliance, the RTS 1000 offers an established Lustre configuration so setup is smooth. Lustre is complex to setup and to administer, but Terascale provides more direct support for Lustre than other vendors do, which is certainly an advantage to us.

"The Terascale solution fits well within the Fulton Center Initiative environment. The performance allows us to keep our compute nodes busy, the pre-defined configuration made initial setup easy, and the simplified approach to adding capacity and throughput will enable us to grow easily as the center grows," said Stanzione.

### **Terascale's Storage Appliance**

The Terascale RTS 1000 Storage Appliance is a Lustre parallel file system based storage appliance offering 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 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.

Terascala leverages Lustre because it is open source, high performance clustered storage 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. Terascala, has an optimized architecture and 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 use of commodity technology 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 Terascala at [www.terascala.com](http://www.terascala.com) or [info@terascala.com](mailto:info@terascala.com) to learn more.

### **About Terascala**

Terascala develops high throughput, high capacity and cost-effective storage solutions. 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, Terascala is based in Avon, MA. Contact Terascala at 508-588-1501 to learn more about how we can solve your storage throughput challenges, or visit our website at [www.terascala.com](http://www.terascala.com)



Terascala, Inc.  
145 Bodwell St.  
Avon, MA 02322  
Tel: 508-588-1501  
Email: [sales@terascala.com](mailto:sales@terascala.com)  
[www.terascala.com](http://www.terascala.com)