



Purdue Rosen Center for Advanced Computing Brings Science to the Art of Infrastructure with BlueArc Storage Solution

“Titan doesn’t bend or break under the strain of starting parallel jobs on hundreds of compute nodes.”

“The high-performance storage we already had in place served pretty well for certain file sizes and types, but did not perform well for others. We’ve moved projects from that system to our BlueArc infrastructure to get consistent, high performance, without any issues for the researchers or my team.”

— Dwight McKay,
director of systems engineering

Summary

To get ahead of steadily-rising demands placed on storage by modeling, simulation and other data-intensive research tasks on campus and off-site, Purdue’s Rosen Center for Advanced Computing has implemented a BlueArc Titan cluster at the heart of its infrastructure. Titan’s performance, capacity and ease of management let the Rosen Center team prepare for the future with an innovative approach to storage services that can match surges in project work as they happen—without overtaxing IT staff.

The Customer

The Rosen Center for Advanced Computing is part of Purdue University’s IT organization, which centrally manages and administers information technology crucial to the university’s day-to-day business functions as well as research and teaching activities. The Rosen Center provides infrastructure resources and services to researchers on the West Lafayette campus and partners nationwide, for data-intensive projects such as climate modeling, nanomaterial simulation and support for the data analysis grid associated with the Large Hadron Collider.

The Challenge

Providing the technology foundation for aggressive research initiatives, the Rosen Center requires a storage infrastructure that responds swiftly to demand for extremely large files from disk and for writing thousands of small files to storage, as well as massive capacity for “scratch” data generated by sophisticated applications. As demand for information technology resources continues to grow on campus and at the university’s satellite locations, Dwight McKay, director of systems engineering, sought to break through storage performance constraints on long-distance computing and extend a successful compute resource-sharing model to storage.

The Solution

The Rosen Center recently expanded its Titan-based storage infrastructure to support not only home directory and application storage, but scratch space as well. The storage system includes Titan 2000- and Titan 3000-series servers and more than 230 terabytes of disk—the majority of which is dedicated to scratch computation files that need not be archived. McKay and team also put the Titan infrastructure at the center of Purdue’s 2008 Supercomputing Bandwidth Challenge entry, to test the system’s ability to respond to demand for huge files over long-distance networks. This helped McKay assess the potential to support a shared-cost model across multiple campuses for unlimited growth without uncontrolled cost.

The Results

Hundreds of scientists, faculty members and graduate students are using the Titan infrastructure to concurrently run highly complex simulations and conduct sophisticated analysis daily. For example, climate modeling researchers run complex mathematical models representing the climate across the entire planet, using dozens of computers at once. These climate simulations produce dozens of gigabytes of output every hour, from each of a few dozen such simulations running at any given time—all of which must be subjected to analysis and eventual archiving. This sort of workload causes the underlying storage system to see constant change, with data continually being created, moved or deleted. The heavy input/output load calls for a storage solution that can handle the influx of data gracefully, with no downtime or degradation of performance.

McKay observes that just one part of the process, staging data for three on-campus researchers, amounts to constant movement of several terabytes in and out of archives. “The nature of research is changing,” he says. “As people increasingly are taking digital data off of remote instruments, using simulations and producing animations to illustrate their findings, they tell us that a stack of eight or 16 nodes in the lab isn’t enough to handle the work.”

With an ever-changing project load, the Rosen Center needed storage capable of writing out small and large files under incredibly demanding conditions—without putting extra pressure on McKay’s team. He explains, “The high-performance storage we already had in place served pretty well for certain file sizes and types, but did not perform well for other workloads or across our larger clusters. We’ve moved projects from that system to our BlueArc infrastructure to get consistent, high performance, without any issues for the researchers or my team.”

But it was Titan’s performance to specification that played a major part in the Rosen Center’s selection. “We benchmarked several storage solutions in house. We ran intensive I/O benchmarks, using more than 200 compute nodes at the same time,” says McKay. “Titan had posted I/O operations to every drive in the rack and was waiting for the responses from the disks. That’s an easy ‘problem’ to solve with more drives.”

Because a storage solution must continue to perform as the infrastructure scales up and out, the Rosen Center is testing the boundaries of supercomputing beyond what researchers require today, with respect to both computation and storage. “Our Steele research cluster—named for a former Rosen Center director—comprises 893 dual-quad-core server machines, with 7144 cores of compute power supporting multiple jobs that each span a section of the cluster,” says McKay. “Titan doesn’t bend or break under the strain of starting parallel jobs on hundreds of compute nodes.”

McKay is even evaluating a “green” supercomputer with capacity for 5,832 cores—and using multiple Ethernet connections to Titan, rather than defaulting to proprietary storage technology. The Rosen Center also put Titan to the test in its the Supercomputing 2008 Bandwidth Challenge entry, designed to demonstrate performance in providing NFS protocol file service over multiple machines and very long distances.

It’s an experiment with a practical purpose. “We are building a statewide grid to provide storage from our West Lafayette campus to our satellite campuses,” says McKay. The work is fundamental to his team’s pursuit of a cutting-edge, standards-based approach to storage. To extract maximum return on investment in equipment and energy use, McKay’s team has pioneered high-throughput computing to harvest idle cycles from any servers that the university owns. Now, McKay believes that he can apply the same concept to storage, with Titan’s performance and capacity as a foundation.

The Conclusion

The “community” model has led to compute resource utilization as high as 80 percent, and McKay anticipates that storage utilization—and researchers’ productivity—will be extremely high as well. The community model includes clusters with data center, racks and networking set up by the Rosen Center, compute cycle sharing and storage made available to research teams and faculty who invest in the compute nodes and storage media they need to support their work. Users will have guaranteed space on the storage infrastructure, and on compute nodes as well as first pick of idle nodes and available scratch storage. In McKay’s view, this approach makes sense not simply to meet researchers’ storage demands in a cost-effective way, but also to be able to run IT as a business rather than as a cost center.



BlueArc Corporation
Corporate Headquarters
50 Rio Robles Drive
San Jose, CA 95134
t 408 576 6600
f 408 576 6601
www.bluearc.com

BlueArc UK Ltd.
European Headquarters
Queensgate House
Cookham Road
Bracknell RG12 1RB, United Kingdom
t +44 (0) 1344 408 200
f +44 (0) 1344 408 202