NetApp
Go further, faster

Success Story

Leibniz Computing Center Achieves 10x Faster Supercomputer Data Services with NetApp

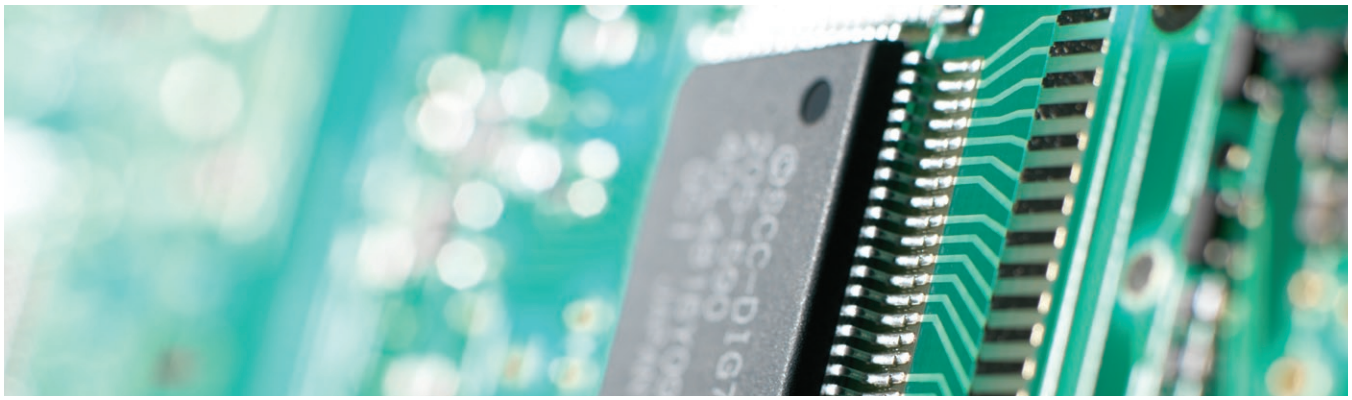

KEY HIGHLIGHTS

Industry
High tech/research–supercomputing

The challenge
Deliver project file services to more than 120 projects in 4,096-CPU SGI Altix 4700 supercomputing environment.

The solution
Six-node NetApp® Data ONTAP® GX system provides high-performance, scalable file services and flexible volume management.

Benefits
- Improved metadata performance 10x
- Enabled fast, transparent capacity allocation
- Integrated replication for highest data protection
- Delivered zero-impact performance disaster recovery
- Increased scalability to double capacity

CUSTOMER PROFILE

Located in Garching, Germany, the Leibniz Computing Center (Leibniz-Rechenzentrum, LRZ) provides general IT services for Munich universities and for the Bavarian Academy of Sciences and Humanities. The LRZ hosts the Munich Scientific Network, a powerful communications infrastructure, and serves as a competence center for data communication networks. Providing large-scale disk and automated tape storage for backup and archiving, the LRZ currently manages more than 1.7 petabytes (PB) of data with an anticipated annual growth rate of 100%. The Leibniz Computing Center for technical and scientific high-performance computing is open to users throughout Germany. (Source: www.lrz.de)

THE CHALLENGE

Ensure fast, reliable delivery of critical project data in a 4,096-CPU supercomputing environment

Few environments exert more demand on a storage architecture than the high-profile supercomputing center at LRZ. The center operates a Linux®-based SGI Altix 4700 with 4,096 Intel® Itanium 2 CPUs, more than 17TB of main memory, and a compute power of 26 TFlops/second. Called HLRB II, the system is one of three national supercomputers in Germany and is ranked 18th on the international "Top 500" list of supercomputers.

The system has been online since July 2006, serving a diverse German technical and scientific community working on projects in physics, material research, fluid dynamics, astrophysics, geosciences, chemistry, life sciences, and medicine.

In such an imposing environment, the storage infrastructure must perfectly complement the compute solution in order for technical and scientific users to achieve desired project results. Dr. Bernd Reiner, a storage specialist at the LRZ, describes the center's data access and availability requirements: "We must be able to handle two different types of project data—project files that include input data, program codes, and final results; and temporary files that contain intermediate results from ongoing simulations.

"The storage requirements for the temporary files and the project results files are very different. Temporary files are often terabyte-sized, but availability and consistency needs are lower. Hundreds of project results files are much smaller in size, but they must be highly available. For the project results, it's essential to have a storage architecture with optimized metadata performance to handle large numbers of relatively small files. Reliability and data consistency are also more critical for project files."

> "With Data ONTAP GX, NetApp delivers an architecture that outperforms comparable systems in terms of performance, consistency, and scalability. We are excited to operate the world's first supercomputing-environment implementation of the NetApp Data ONTAP GX system."
>
> **Dr. Herbert Huber**
> Head of Compute Server Group, Leibniz Computing Center

Reiner points out that because the LRZ routinely manages a changing mix of more than 100 projects, essential storage solution elements include manageability, reliable backup, and zero-impact performance disaster recovery. "End to end in the project lifecycle, our users depend on us to deliver the best possible working conditions and to ensure that their project data is protected and accessible."

THE SOLUTION

Consolidate and standardize on NetApp Data ONTAP GX for high-performance computing file services

To evaluate storage architectures for project file services in its supercomputing environment, the LRZ developed a series of benchmarks. Based on real-world project requirements and conditions, the tests evaluated both metadata performance and sequential data throughput. Taken together, the values indicate system performance when a file with output data is stored in a project share.

Christoph Biardzki, a storage expert at the LRZ, says that strong benchmark results led to the implementation of the NetApp Data ONTAP GX system. "We had been using NetApp solutions for several years for file services, data protection, and archiving, so the management benefits and high reliability of NetApp's architecture were well known. What we needed to test was the performance of the Data ONTAP GX system in the HLRB II environment. We found that NetApp's architecture outperformed comparable systems, and we moved forward, putting the world's first Data ONTAP GX system installation into production."

The NetApp Data ONTAP GX system is deployed in a scale-out configuration based on NetApp FAS3050 building blocks. Currently, the six-node system provides 40TB net capacity on Fibre Channel disks connected over NFS to a total of 16 supercomputer partitions with 256 processors each. The NetApp system provides storage for all HPC project data and software. The Data ONTAP GX system's global namespace feature allows several independent data storage systems to be presented and managed as a single large storage pool containing all files and file systems.

"With Data ONTAP GX, NetApp delivers an architecture that outperforms comparable systems in terms of performance, consistency, and scalability.

We are excited to operate the world's first supercomputing-environment implementation of the NetApp Data ONTAP GX system," stated Dr. Herbert Huber, head of the Compute Server Group at Leibniz Computing Center.

The LRZ leverages NetApp Snapshot™ technology to enable online protection of the critical project data stored on the Data ONTAP GX system. Integrated SATA disks providing 10TB of storage provide economical replication within the cluster. Backups are also written via NDMP onto STK T10000 tape drives. For long-term archiving, users can access an archive system with a 2,300TB-capacity tape robot environment managed by IBM Tivoli Storage Management (TSM) software.

BUSINESS BENEFITS

Boosted performance 10x while achieving data protection, seamless scale-out, and simplified management

"With a synthesis of global namespace functionality and high performance plus reliability, scalability, and manageability, NetApp's solution is the ideal storage architecture for our supercomputing environment," comments Biardzki.

"NetApp's solution is a perfect match for applications that handle numerous files, need a high transaction speed, and require absolute data integrity," adds Reiner. "Even the implementation was faster and easier than we expected. Throughout the process, we were supported by local NetApp engineers, and a specialist from NetApp Israel

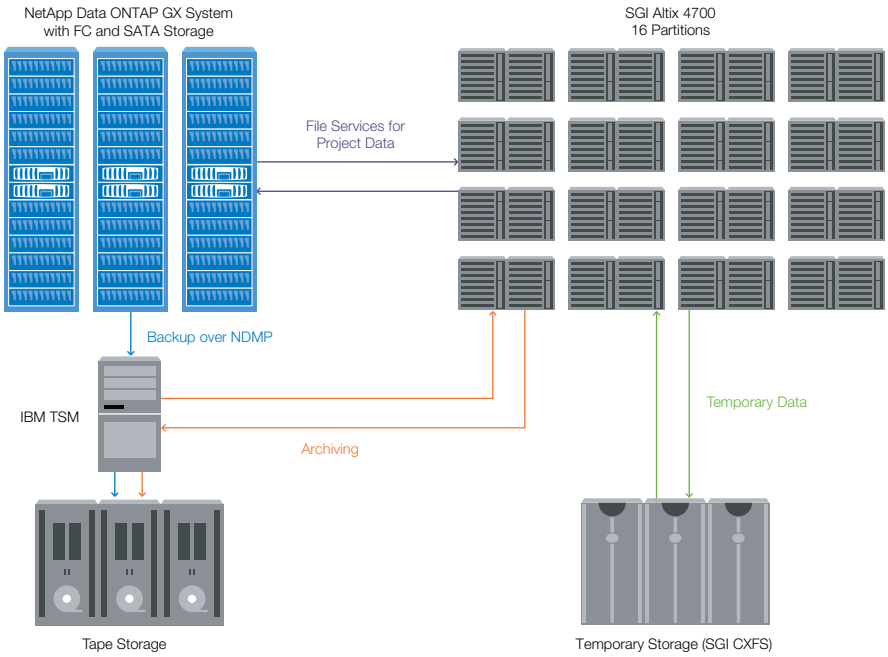

Figure 1) The storage structure of HLRB II at LRZ.
The file services, data processing, and output storage implementations are part of the complete data security solution, which includes online Snapshot copies and backup and archiving on high-speed StorageTek tape robots, supported by disk caching. IBM TSM controls the archives, to which users have direct access.

helped with the fine-tuning. Based on our experience, we would definitely choose NetApp again."

10x performance boost
Biardzki reports that day-to-day project operations at the LRZ have proven the performance of the Data ONTAP GX system. "The performance improvements have been noticeable to our users. Besides high throughput, which relies more on physical characteristics, our project-shares environment requires consistently high metadata performance. Metadata handling had been a problem with our previous supercomputer, the HLRB I, a 1,368-CPU Hitachi SR 8000 with direct-attached disk storage. In that system, the metadata was stored separately from the data, so performance degraded with increasing requests for files."

In contrast, NetApp's system maintains data and metadata in the same volume, optimizing and ensuring consistent high performance for environments that require random access to hundreds of small files. "Compared to the HLRB I, the metadata performance of HLRB II with NetApp storage is 10 times greater," points out Biardzki. "The NetApp architecture makes it the ideal storage solution for our project shares."

Streamlined project support, management, and scalability
Today, dedicated capacity on the NetApp storage is allocated to each scientific project with compute allowance on the HLRB II. A national committee of renowned scientists presides over the project approval process as well as the determination of compute-time contingents. In order to receive approval, applicants must have experience working with supercomputers and parallel processing, and the project must fit the computing architecture of the HLRB II. Once a project is approved, LRZ sets up the compute schedule, provisions project shares, and validates the project on the HLRB II.

Biardzki says that the Data ONTAP GX system's global namespace functionality streamlines storage management for the complex project mix supported on the HLRB II supercomputer. "In traditional architectures, file servers hit performance or capacity limits, forcing the addition of new file servers. Because the systems in such an environment operate independently, they have to be managed separately, and users must typically spend a lot of time localizing individual files.

"The global namespace dramatically simplifies the process of capacity and performance scaling by effectively allowing us to summarize all of our files and file systems under a common namespace, even though the files physically reside on different storage systems. In our environment, each project has its own storage space that is integrated into the global namespace and can be moved as necessary within the cluster. This administrative functionality was very important for us as capacity adjustments and reconfigurations must be done on the fly to support our project demands. This architecture also ensures that data entry, simulation computing, and data output processes operate independently from storage management tasks."

With more than 120 parallel projects to manage, the LRZ storage team benefits significantly from the NetApp global namespace and flexible volume management capabilities. Biardzki adds, "Our projects vary dramatically in terms of data set size, number of users, shared data-access requirements, and storage needs, so we have to be able to quickly and easily allocate, change, or scale out capacity. The global namespace makes the physical storage layout totally transparent to the users—all they see are their own project spaces.

> "NetApp delivers high performance and high availability with the convenience of network-attached storage. We would choose NetApp again. The NetApp Data ONTAP GX system is a perfect match for applications that handle large numbers of files, need a high transaction speed, and require absolute data integrity."

Dr. Bernd Reiner
Storage Expert, Leibniz Computing Center

"NetApp's solution also offers performance and capacity scalability to support future expansion of the LRZ compute infrastructure. As we scale the throughput and capacity of the NetApp system, we will be able to maintain consistent metadata performance. And, because of the global namespace, we'll be able to accomplish the logical integration with just a few mouse clicks."

Consistency and data protection

Biardzki says that the implementation of a storage platform dedicated to project file services was critical to ensuring data consistency and protection. "We needed an architecture that could support our performance requirements while guaranteeing the quality and correctness of compute results. With NetApp's solution, we can be sure that data is filed and recalled correctly. Project shares are fully protected with integrated Snapshot copies and replication—NetApp's solution gives us a very convenient way of doing online backup, mirroring, and disaster recovery. Other file systems couldn't offer equivalent functionality at all or could only deliver it with a severe performance penalty. The NetApp Data ONTAP GX system absolutely delivers the best storage feature set and performance for our supercomputing project shares."

"NetApp delivers high performance and high availability with the convenience of network-attached storage. We would choose NetApp again. The NetApp Data

ONTAP GX system is a perfect match for applications that handle large numbers of files, need a high transaction speed, and require absolute data integrity," concludes Reiner.

SOLUTION COMPONENTS

NetApp Products

NetApp Data ONTAP GX

FAS3050 with FC and SATA disks

NetApp Snapshot technology

Third-Party Solutions

IBM Tivoli Storage Management software

NetApp creates innovative storage and data management solutions that accelerate business breakthroughs and deliver outstanding cost efficiency. Discover our passion for helping companies around the world go further, faster at www.netapp.com.

NetApp™

www.netapp.com