

CASE STUDY

EFFICIENT OPERATIONS AND FAST DATA ACCESS ENABLE BOTTLENECK-FREE RESEARCH AT HLRS

Quobyte software storage accelerates engineering and scientific breakthroughs by simplifying storage management, delivering HPC-required parallel file system performance, and facilitating data ingest and processing.



Summary

Needs/Challenges: For an infrastructure overhaul, the HPC experts at HLRS were looking for a storage system that could fulfill three major requirements – apart from performance: For one, it had to be scalable to cope with future data growth; for another, it must be manageable without increasing the number of staff; and finally, it needed to provide S3 object storage capabilities to enable easy access for researchers and teams using the infrastructure.

Solution: Quobyte's Data Center File System

Platform: + 8 high-density servers with a 60-disk JBOD each
+ Extreme Networks Summit x870 switch fabric
+ 25 GbE network infrastructure
+ NEC QxFS with NEC support
+ A hybrid of a parallel file system and S3 storage

Use Case: Commercial and scientific HPC

Key Benefits

- High-performance parallel file system with native S3 access for faster data ingestion and processing
- Automated management enables low-touch operations with only a small team of admins
- Eliminates storage silos and enables collaborative research thanks to truly unified storage
- Multi-tenancy provides strong separation of concern
- Overall gains in operational efficiency

“ The two crucial features – performance and scalability – are just part of what makes Quobyte the perfect storage system for HLRS. It's the software's multi-tenancy capabilities, fault tolerance, and manageability that are lifting it way above par. ”

Dr. Thomas Bönisch
Co-Chair of the HLRS innovation group

HLRS – Extending the Boundaries in Scientific Research and Commercial HPC

The High-Performance Computing Center Stuttgart (HLRS) is running one of the world's top supercomputing sites. In close partnership with NEC Deutschland GmbH, they run additional high-performance clusters focused on providing efficient compute and storage resources for data-intensive workloads.

Founded in 1995 under the auspices of the University of Stuttgart, the HLRS is one of three members of the Gauss Center for Supercomputing (GCS), an alliance of Germany's top three supercomputing sites. The HLRS provides scientific researchers and industry organizations with high-performance computing platforms and technologies, including services and support. The most demanding workloads in various industries and research sectors run on HLRS' infrastructure, ranging from engineering and health to climate and mobility. The millionth job running on the supercomputer Hazel Hen, for example, was focused on multiphase flows in non-Newtonian fluids – things like toothpaste or paint – which behave differently from what's expected assuming Newton's laws of viscosity. Running virtual experiments, the researchers learned about similarities in flows for both kinds of fluids, enabling engineers to improve the efficiency of nozzle Designs.

In addition, the HLRS conduct basic and applied research in high-performance computing themselves, always improving upon the efficiency of large-scale data workloads.

H L R I S

High-Performance Computing Center | Stuttgart

Meeting the Demands of Capacity & Performance

Running a supercomputing site is not just about racking up compute power, but storage is also a chief concern. No matter the workload – be it automotive engineering or climate simulations – fast and reliable data access is paramount in HPC. While scaling up compute power is mostly a matter of cost, avoiding storage bottlenecks is more about smart software these days.

When HLRS entered into the next planning for an infrastructure update, storage was a major focus and had to meet three essential requirements apart from delivering performance: First, it had to be scalable to secure future growth. Second, and related to the first point, it had to be manageable at scale so that capacity growth could be managed by the same team. Third, an object storage solution was needed to simplify management of data coming into and going out of the system or, in other words, to make data access more convenient for research teams using the infrastructure.

HLRS runs very demanding workloads – for industry and research organizations alike. Whether they're testing the crash behavior of cars or simulating turbine flows, aerodynamics, or the load-bearing capacity of structural elements, the engineering-oriented enterprises and institutions require high-performance storage access for an exploding amount of data. Hence, scalability and performance are the two indispensable features for HLRS' storage infrastructure.

The teams of HLRS and NEC were jointly looking for a solution that fulfilled both these major needs, performance and scalability – and then some. For NEC, the choice was clear: they had known and worked with Quobyte's founders since before the company was founded, during the early days of XtremFS.

Thanks to his intimate knowledge of Quobyte's Data Center File System, NEC's senior R&D engineer Dr. Erich Focht knew right away that it was the optimal solution to HLRS' setup: "I've experienced what great performance Quobyte's parallel file system delivers and I know that the software was designed with scalability as a primary concern. And the importance of scalability can't be overstated for HLRS' HPC environment: engineering workloads or climate simulations aren't just compute- but also, even more so, data-intensive."

Enabling Collaboration by Removing Data Silos

"The two crucial features – performance and scalability – are just part of what makes Quobyte the perfect storage system for HLRS, though," says Dr. Thomas Bönisch, Co-Chair of the HLRS innovation group, "it's the software's multi-tenancy capabilities, fault tolerance, and manageability that are lifting it way above par."

Unlike other HPC file systems, Quobyte delivers features that prove extremely useful for modern HPC infrastructure management. Multi-tenancy with full-fledged ACLs means that a strict separation of concerns and users is possible, guaranteeing a high level of security. At the same time, the ability to

host different users on one and the same file system pleases the HLRS administrators because it makes the entire system more manageable. Instead of having to manage several different storage systems, they deal with a single interface that serves all their needs. Hence, there are no more data silos that would require additional manpower to handle and operate.

"The hybrid of S3 and parallel file system saves both researchers and administrators valuable time and makes their jobs a lot easier!"

— Thomas Beisel, Head of Software & Systems, HLRS

Easy Access on all Interfaces

What sets Quobyte apart even more is a unique feature: its native S3 interface. This allows researchers to easily ingest data via the standard object storage protocol and run analytics on that same data right away via the parallel file system. Creating a single namespace for all of HLRS' data assets eliminates the need for tedious data migrations within the storage system. "The hybrid of S3 and parallel file system saves both researchers and administrators valuable time and makes their jobs a lot easier," comments Thomas Beisel, Head of Software & Systems at HLRS.

Beyond S3, Quobyte supports a wide range of protocols that make it HPC-suitable – it integrates with Hadoop, speaks NFS and SMB via Samba, and has native clients for all OS types: Windows, MacOS, and Linux.

Performance and Manageability at Scale

The operational efficiency HLRS' team gains through the unified storage approach – that it grants access to the same data from virtually any platform – is further flanked by Quobyte's patent-pending Paxos implementation, which allows the system to smartly handle hardware failures on its own; it automatically routes around such failures and avoids service interruptions. "The built-in monitoring via GUI or API helps administrators get a quick and detailed overview of what's happening with the storage resources and where," says Dr. Erich Focht. Scaling out by adding further devices and exchanging faulty hardware becomes a breeze.

All in all, Quobyte's Data Center File System is the perfect match for HLRS' HPC storage requirements. It scales to hundreds of PBs without increasing the administrative burden and handles workloads of hundreds of millions of small files with ease. Multi-tenancy and ACLs enable secure and strict separation of concerns so that different research teams and companies can use and work on the same infrastructure simultaneously. And Quobyte's monitoring tools, together with the health manager and a clear dashboard provide a near real-time overview of the cluster's status and health, providing storage administrators the tools they need to drill into the details required to operate efficiently.



Quobyte®

Data Center File System.™ Fast and Reliable Software Storage.

Quobyte Inc.
4633 Old Ironsides Drive
Santa Clara, CA 95054
info@quobyte.com
+1 650-564-3111

Quobyte GmbH.
Hardenbergplatz 2
10623 Berlin
info@quobyte.com
+49 30-814 591 800