



CASE STUDY

GWDC Enables Breakthrough Research with StorNext at University of Göttingen and the Max Planck Society

GWDC—university computing center for the University of Göttingen and computing and IT competence center for the Max Planck Society—provides IT resources and consulting services to scientists and other researchers from the university and the Max Planck Institutes. The Quantum StorNext® multi-tier platform helps GWDC store and provide access to huge volumes of valuable data.



FEATURED PRODUCTS



“In our experience, StorNext is the only platform that can handle the number of files and amount of data we manage while also delivering high performance.”

Stefan Teusch
Deputy head of IT infrastructure, GWDC



StorNext has met our needs even as our workflows have become more demanding, our environment has evolved, and our management team has changed. We're confident that StorNext will continue to meet our data management needs well into the future.

Dr. Philipp Wieder - Deputy head, GWDG



SOLUTION OVERVIEW

- StorNext Scale-out Storage, including
 - StorNext Metadata Appliances
 - StorNext Gateway Appliances
 - StorNext AEL6000 Tape Archives

KEY BENEFITS

- **Manages billions of files and petascale data** across multiple tiers, in a single file system
- **Scales to support growing data volumes** that are nearly doubling every year
- **Enables high-speed HPC processing** via the StorNext DLC protocol over IP
- **Provides heterogeneous access to data** across a full range of operating systems
- **Protects data and minimizes downtime** with backup and disaster recovery integration
- **Allows a small staff to manage a big archive** by reducing administrative complexity

GWGDG—university computing center for the University of Göttingen and computing and IT competence center for the Max Planck Society—enables university researchers and scientists from the prestigious Max Planck Society to conduct groundbreaking studies and to support innovative initiatives across a broad range of subjects. For example, while biomedical researchers capitalize on the center's high-performance computing (HPC) clusters to analyze genomic data, humanities groups are digitizing files to make library content more easily accessible.

FACING INCREASED DEMAND FOR DATA STORAGE AT UNIVERSITY OF GÖTTINGEN

GWGDG's IT group needed a storage environment that could scale to meet rapidly rising demand. "The amount of data we store nearly doubles every year, and the growth rate continues to increase," says Dr. Philipp Wieder, deputy head at GWGDG.

Data growth is driven partly by the quest to answer more complex scientific questions. "Next-generation sequencing, for example,

produces huge amounts of data," says Wieder. "One team needed to store 300TB of data just for one project."

GWGDG also provides IT resources for groups outside of science. For example, the computing center supports the Göttingen State and University Library—which serves the university, the German State of Lower Saxony, and the Göttingen Academy of Sciences. "The library is digitizing texts and creating models of physical objects," says Wieder. "All that data must be stored somewhere."

The computing center needed to manage the large archive and make data available to users—today there are more than 40,000 users. "We needed a file system that could handle billions of files and tremendous data volumes," says Stefan Teusch, deputy head of IT infrastructure, GWGDG.

CREATING A SINGLE MULTI-TIER STORAGE ENVIRONMENT THAT SPANS DISK AND TAPE

The storage platform also needed to support a multi-tier hierarchical storage management (HSM) environment. "We now have in StorNext

2.5PB of data on disk and 4.5PB on tape,” says Teusch. “The goal was to make the data accessible from a single file system.”

The economics of tape libraries—including their low power consumption—helps the center meet the requirement to store data for ten years or more. “Tape is more economical than disk,” says Teusch. “And since much data is accessed infrequently, it makes sense to use tape.”

STORING BILLIONS OF FILES AND 7PB OF DATA WITH STORNEXT MULTI-TIER STORAGE

To manage valuable data assets, the center adopted the Quantum StorNext platform more than a decade ago. The center has expanded and adapted the storage environment easily over the years. Today, the IT group manages 7PB of data with a StorNext multi-tier storage solution, powered by StorNext data management software. The solution includes primary disk and StorNext tape archives.

“In our experience, StorNext is the only platform that can handle the number of files and amount of data we manage while also delivering high performance,” says Teusch.

EASILY SUPPORTING ALL THE DIFFERENT OPERATING SYSTEMS THAT SCIENTISTS USE

The StorNext platform enables the center’s IT group to support users across a full range of operating systems. “Many scientists, particularly in the natural sciences, use UNIX or Linux, but other users might prefer Windows or Mac OS. In fact, some departments are Mac-only,” says Teusch. “With StorNext, we can easily support all these operating systems. Users can access the data they need without making any changes to their workflows.”

ENABLING HPC PROCESSING WITH HIGH-SPEED SHARED STORAGE OVER ETHERNET

GWGD is a well-respected HPC provider in Germany. To provide the center’s HPC cluster with shared access to data from the StorNext environment, the IT group installed StorNext distributed LAN clients (DLCs) on HPC nodes—and deployed seven StorNext gateways for load balancing and failover. “The StorNext DLC

eliminates the need to connect all HPC nodes with Fibre Channel,” says Teusch. “We can continue to use traditional Ethernet.”

The StorNext DLC protocol provides high-speed access to data over Ethernet to enable the HPC processing. “HPC users requested 2GB-per-second performance. With StorNext DLC, we’re achieving 3.2GB per second,” says Teusch. “Researchers can quickly move large amounts of data into the HPC cluster for processing and out again for long-term storage.”

PRESERVING COPIES OF DATA FOR DISASTER RECOVERY

To protect online data, the IT group uses StorNext policy-driven tiering to create copies of data residing in the primary disk cache. “If a hardware device fails or we lose a storage system, we still have the copy in the StorNext tape archive, and we can rapidly restore data,” says Teusch. “Users can get right back to work.”

MANAGING A LARGE AND GROWING SCIENTIFIC ARCHIVE WITH A SMALL TEAM

Even as the multi-tier storage environment continues to grow, managing the StorNext platform has remained simple. The IT group can continue to expand without hiring extra staff or pulling existing staff away from other projects. “We invest feasible and predictable effort to manage our StorNext environment—it is very easy,” says Wieder.

ENVISIONING AN INTEGRATED NETWORK OF STORNEXT ENVIRONMENTS ACROSS GERMANY

Looking ahead, the center’s IT group is considering new ways to employ StorNext. “Some organizations that use our StorNext archive, such as the Max Planck Society, have their own StorNext deployments. They keep their primary data in their local StorNext system and use our environment to archive a second copy,” says Wieder. “We’re considering ways to facilitate similar integration with organizations that are farther away. We could build a network of connected StorNext environments to help streamline collaboration across Germany.”

“We invest feasible and predictable effort to manage our StorNext environment—it is very easy.”


Dr. Philipp Wieder,
Deputy head,
GWGD

Left and right coiling snail and two frog embryos.
Photography by Esther de Roij and Gary McDowell;
digital composition by Jeremy Guay, Peregrine Creative.



ABOUT GWGD

Gesellschaft für wissenschaftliche Datenverarbeitung mbH Göttingen (GWGD) is a computer center shared by Germany’s Göttingen University and the Max Planck Society. GWGD fulfills the functions of a joint computing and IT competence center, providing a range of basic IT services, data management solutions, and computing facilities to support scientific research and promote professional training. The center has built a large-scale storage environment to store, protect, and provide access to millions of files and large-scale data volumes for tens of thousands of users.




Real-time measurement of the start of an ultrashort pulse laser.
Illustration: Georg Herink, University of Göttingen.



EASILY HANDLING NEW REQUIREMENTS AND DATA GROWTH WITH QUANTUM STORNEXT

The GWDG IT group sees StorNext as an ideal storage solution that meets the center's needs for scale, performance, and manageability. "StorNext has met our needs even as our workflows have become more demanding, our environment has evolved, and our management team has changed," says Wieder. "We're confident that StorNext will continue to meet our data management needs well into the future."



Quantum StorNext scale-out storage enables the GWDG to manage a multi-tier, petascale environment of valuable data.

To contact your local sales office, please visit www.quantum.com

Quantum