

Case Study

Siemens Healthineers Digital Technology and Innovation


See how Siemens Healthineers manages storage for its AI Factory infrastructure using Quobyte

Digital Technology and Innovation (DTI) is the central hub for R&D in artificial intelligence and digital innovation of Siemens Healthineers. DTI's global footprint extends to USA (Princeton, NJ), China, India, and various countries within Europe. DTI's highly skilled AI experts specialize in using large data collections and powerful supercomputing infrastructure to build artificial intelligence solutions in close collaboration with government agencies, universities, and healthcare providers worldwide.



DTI's supercomputing infrastructure is built on state-of-the-art high-performance computing equipment for Deep Learning on Big Data. This includes a supercomputer named "Sherlock", which consists of a cluster of Linux-based HGX-H100 and HGX-A100 compute nodes with NVIDIA GPUs totaling 340 Petaflops, tens of Petabytes of storage including 6 Petabytes all-flash NVMe, 200 GBps InfiniBand networking,

10 Gbps fiber links to the cloud, and 10 Gbps point-to-point fiber connectivity to two additional edge supercomputers in Germany and China. The main datacenter in the U.S. is in Edison, NJ. Sherlock's GWh-scale power consumption relies 100% on renewable energy from solar and wind, in line with Siemens Healthineers' sustainability goals.



Gianluca Paladini, Sr. Director of Engineering, heads the AI Supercomputing group in charge of such global HPC infrastructure. "Our scientists use Sherlock to run an average of 1,600 deep learning experiments per day, training AI models using a Data Lake with billions of medical images, medical reports, lab tests, genomic data, and treatment plans," said Paladini. "The AI Supercomputing team provides the processing power, software stack and HPC expertise to parallelize and accelerate training of AI models produced by DTI's AI Factory."

The AI Factory, established at DTI since 2016, consists of processes, resources and infrastructure for data collection and curation, data governance, structured annotation, model training, testing, prototyping, and clinical validation studies in order to conduct the necessary research to develop AI models that perform with high accuracy, and deliver them to Siemens Healthineers business lines for future integration into products. DTI has delivered hundreds of AI models, which as of May 2024 have been integrated into 60+ FDA approved AI-enabled medical devices¹.

¹U.S. Food and Drug Administration, Artificial Intelligence and Machine Learning (AI/ML)-Enabled Medical Devices, (web link).

The Challenge

Since its inception in 2016, the Sherlock supercomputer has undergone upgrades to double its processing speed every year. This exponential growth was driven by the need to run a larger number of deep learning experiments and using much larger training datasets from a rapidly growing curated Data Lake. As storage requirements grew larger, the AI Supercomputing team needed to expand the infrastructure and transition from large-capacity individual NFS servers to a scale-out parallel file system.

After testing Lustre on top of ZFS and subsequently BeeGFS, the team tried Quobyte and was impressed with how easily it could transparently recover from hardware failures both at the drive level as well as at the server level. “We used other open-source parallel file systems mainly as scratch space since we did not have the confidence to roll them out as part of a production environment,” said Ron Trickey, Sr. HPC System Administrator. “But with Quobyte, it became clear that we could sustain significant hardware losses with no interruption.”

As AI training datasets grew larger, the approach of copying data to scratch space prior to running HPC workloads was taking too long. “Data size for each deep learning job can be several Terabytes and copying it to scratch space is impractical. I needed to find a production-ready scalable storage solution that could be accessed directly by AI training workloads at high speed, without sacrificing reliability, maintainability, and fault-tolerance,” said Paladini.



Data size for each deep learning job can be several Terabytes and copying it to scratch space is impractical. I needed to find a production-ready scalable storage solutions that could be accessed directly by AI training workloads at high speed without sacrificing reliability , maintainability, and fault-tolerance



Gianluca Paladini
Sr. Director Of Engineering,
AI Supercomputing

Because the AI Supercomputing group had already made a considerable capital investment in numerous storage servers with tens of Petabytes of capacity, including several Petabytes of all-flash NVMe, the ideal system would have to run on our existing hardware. Many storage vendors were proposing solutions that required software to be bundled with new hardware, which would have been too costly.

The Solution

- ✓ **Low Operating Costs** thanks to easy configuration, automatic tiering, and the use of commodity hardware
- ✓ **Fault tolerance** with 3-way replication and erasure coding, with no downtime during loss of disk drives or entire storage servers
- ✓ **Unified Storage System** supporting heterogeneous data access and advanced security features for data governance
- ✓ **Scalable Capacity and Performance**, with I/O throughput that scales up linearly as capacity increases with more storage nodes

When looking at Quobyte's extensive feature set, it became clear that it checked all the boxes to fulfill our requirements of a production environment to be used for HPC workloads. In addition to providing high-performance and fault-tolerance, Quobyte would provide advanced security features for our Data Governance processes, facilitate web-based data collection from clinical collaboration sites, and support with a variety of file access protocols compatible with Linux and Windows.

The migration to Quobyte was carried out gradually starting in 2021, without disrupting production work. "Since it runs on commodity hardware, we were able to migrate data and storage capacity to the Quobyte cluster one server at a time, reusing all the storage hardware we had already purchased. Quobyte rebalanced all the data across the cluster automatically," said Paladini. "The whole process was seamless to users – after each server was migrated, we would simply redirect their access to a different NFS mount of the same name stored on Quobyte."

The Results

Quobyte as the ideal parallel file system for an AI Factory

With Quobyte, the AI Supercomputing team has finally achieved the desired AI Factory scalability and operational efficiency. Simplified configuration management reduces maintenance overhead, and operating costs are substantially lower. The Policy Engine is very efficient at file placement, it can be easily configured to automatically determine which files should reside on fast NVMe drives used by HPC workloads, and which files should be archived onto slower and cheaper HDD storage, without having to provision separate physical storage tiers. Quobyte's unified file systems supports SAMBA, facilitating the transfer of project files from the scientists' Windows-based laptops. POSIX compatibility makes it easier to use all the latest and greatest Python-based AI tools on the Linux-based supercomputer; and Quobyte's compatibility with Spark/HDFS makes it easier to process non-imaging sparse tabular data.

More importantly, Quobyte makes it very easy to add more storage capacity without requiring any outage or downtime. "Storage requirements increase considerably for every active research project, because data preparation for AI model training involves data preprocessing, augmentation and annotation steps that generate millions of additional files," said Pragnesh Patel, HPC Software Architect in the AI Supercomputing team, "thanks to Quobyte's scalable high-throughput, Sherlock has been able to carry out very large deep learning experiments which are at the forefront of AI research."

For example, in 2022 large-scale self-supervised training experiments were conducted on datasets consisting of over 100 million medical images². "More recently, Sherlock was used to train large deep neural networks using self-supervised learning from more than half-a-billion images, requiring massive I/O," said Paladini. "Inspired by how humans learn a little from everything, this approach can learn from any imaging modality, with or without annotation, parsing tirelessly through 500 million images – a process that would take a radiologist over 150 years, assuming 10 seconds per image with no sleep and no break. And this trend is not slowing down – Sherlock is now training Large Language Models with billions of parameters, requiring massive storage for training datasets."

²"Contrastive self-supervised learning from 100 million medical images with optional supervision", Journal of Medical Imaging, Vol. 9, Issue 6, 064503, November 2022 (web link)

Improved Data Governance, Data Privacy and Security

All Data Governance aspects of the AI Factory process are driven by a team called Big Data Office (BDO). Such process is greatly facilitated by leveraging key features of the Quobyte parallel file system provided on Sherlock by the AI Supercomputing group.

BDO collects de-identified data from a vast network of clinical collaboration sites. Quobyte's compatibility with the S3 API eliminated the need to upload data to the cloud first, and then transferring it to our on-prem supercomputer. Such cloud copy was time consuming and costly for Petabytes of data. Instead, the AI Supercomputing team setup an easy-to-use

secure file transfer service based on GoAnywhere, which interfaces with Quobyte's S3 API so that data can be received and stored directly and securely, without an intermediate copy on a public cloud.

After receiving data collections from clinical collaborators, BDO performs a series of curation steps in order to verify the quality and integrity of both data and its associated metadata. Quobyte's Tenant Domains feature plays an important role during this phase, as it guarantees that data can be quarantined in a tenant space completely isolated from the rest of the file system, until BDO clears it and releases it into a Data Lake for use in R&D projects.

The Data Lake requires stringent data access controls since each data collection has different data access and retention policies. It takes advantage of Quobyte's support for unified Access Control Lists (ACLs) which remain consistent when accessing files in Linux and Windows, so that the BDO team can enforce which users are authorized to access certain files for a specific R&D project. For increased security Quobyte also provides automatic end-to-end AES encryption.

Comprehensive Fault-Tolerance for Business Continuity

With over 2 billion curated medical images and billions of data points collected and managed by BDO's team, Data Lake storage is mission critical for

more than 150 AI Factory R&D projects/year, requiring several Petabytes for Data Lake collections and additional storage for active research projects. It is therefore imperative to setup fault-tolerant infrastructure that can provide business continuity for AI Factory operations.

Running the production environment with Quobyte's fast 3-way replication ensures the ability to survive occasional disk failures seamlessly with no downtime, and even the simultaneous failure of two entire storage servers without any data loss. But for true business continuity, the AI Supercomputing team also provisioned a separate secondary cluster which uses Quobyte's multi-cluster asynchronous replication capability. The secondary cluster is within the datacenter on the same high-speed network – this way in the eventuality of a catastrophic failure or a long maintenance outage, the supercomputer can immediately switch to the secondary mirror cluster and continue to work uninterrupted at high bandwidth.

In addition, the datacenter is connected via dedicated point-to-point fiber to a disaster recovery location where a backup copy is stored. These DR clusters leverage Quobyte's Erasure Coding feature, which stores data more compactly than 3-way replication, thereby providing additional cost savings. With such redundancy we can ensure valuable AI training data is properly safeguarded and business continuity is achieved.



About Quobyte

Quobyte provides a unified storage platform engineered for simplicity and high performance, suitable for enterprises and researchers pushing the frontiers of innovation. It ensures exceptional reliability and enterprise-level features, enabling rapid deployment and streamlined operations for organizations of various sizes, making HPC storage accessible for all.

Quobyte empowers small teams to run large-scale infrastructures, advancing today's most groundbreaking research.

[Learn More](#)

