BeeGFS®

Case Study

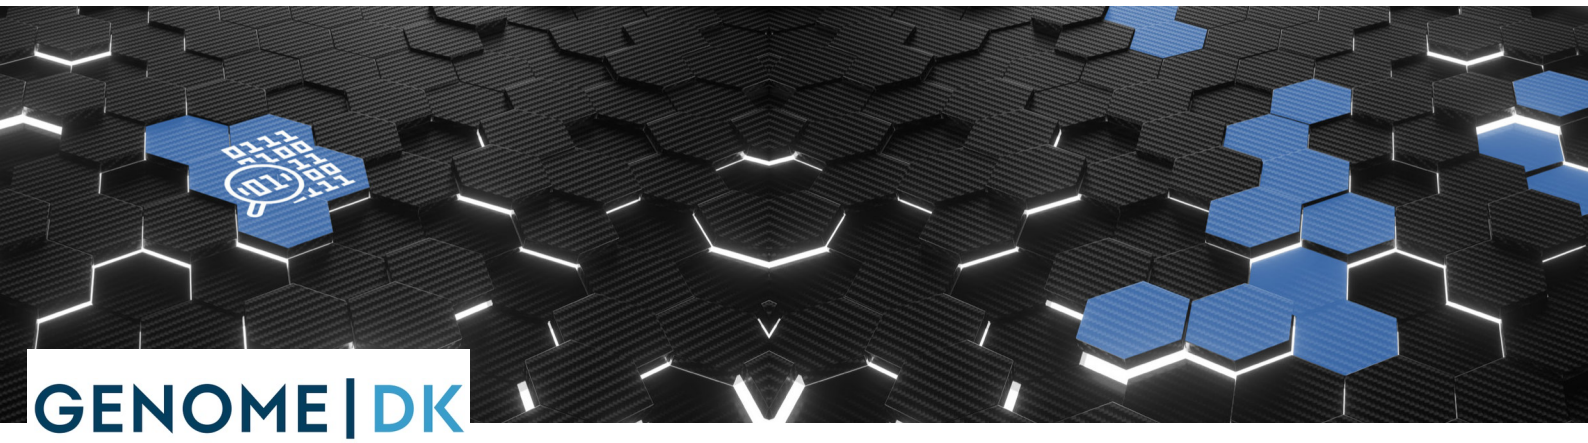

GENOME | DK
HIGH-PERFORMANCE COMPUTING

33 Petabytes, One Simple Mount:

Inside GenomeDK, BeeGFS-Powered Research Supercomputing Center.

33 PB
Data under Management

~1300
Active researchers and SMEs

~5000
Concurrent Jobs

Introduction

Founded around 15 years ago by two university researchers who secured grants for research projects involving very large, and sometimes sensitive, datasets. At the time, Denmark had no existing facility equipped to analyze data at that scale, so the two researchers pooled their resources and built their own.

Today GenomeDK exists to support researchers, students, SMEs, and hospitals with the infrastructure to store, manage, and analyze large-scale datasets, as well as run large-scale simulations. GenomeDK aims to be an “all purpose” HPC facility and workloads are therefore extremely varied and change by the minute. This, and the fact that GenomeDK handles clinical workloads for Denmark’s second largest hospital, means that reliability and operational simplicity is key.

The center is publicly funded, and around 200 researchers actively use the system at any given time; with projects routinely including international collaborators. A small team of just three and a half system administrators runs the entire operation end-to-end - hardware, software, tooling, security, compliance, support, documentation, training, and day-to-day operations.

“BeeGFS is just very simple to understand, operate, and troubleshoot.”

Dan Søndergaard, HPC senior consultant, GenomeDK

3.5
System Admins

~12 years
Running BeeGFS

Challenge

GenomeDK's largest storage system is a 33 PB BeeGFS deployment used as a single mount point for everything, active analysis as well as long-term retention. Unlike a typical HPC environment where a fast parallel file system is reserved purely for scratch space, GenomeDK treats the same BeeGFS instance as both working storage and permanent archive; some datasets on the system are 15 years old. On a fixed annual public-sector budget, the team can't simply buy one platform and run it unchanged for five years, instead, storage is expanded roughly every two to three years, with data migrated onto new capacity and the oldest hardware retired, an ongoing, incremental process rather than a single large refresh cycle. This process is performed in the background with minimal impact on services and is aided by internally developed tooling.

Before BeeGFS, GenomeDK ran an enterprise platform built on Panasas that, according to the team, simply didn't hold up in operation. The team evaluated Lustre and BeeGFS as replacements; BeeGFS won out because it was significantly easier to install and operate. A decision made roughly 11–12 years ago, making GenomeDK one of the earliest BeeGFS deployments.

Why BeeGFS

Operational simplicity, at every layer

Many storage systems relied on kernel modules for all components that had to be loaded and unloaded, making the stack feel fragile. BeeGFS, which runs largely in user space, avoids that fragility - the team describes it as fundamentally "a way of collecting a lot of disks into one big folder," a model that's easy to reason about and easy to troubleshoot when something goes wrong.

Built-in visibility into a chaotic, shared workload

With ~200 researchers and up to 5,000 jobs running at once, a single misbehaving job, often one performing an unusually high volume of small metadata operations against a single file, can slow the whole file system down, and it's rarely the job consuming the most compute that's responsible. BeeGFS's observability tools give the team a fast-starting point for tracking down the culprit.

No hardware vendor lock-in - a requirement, not a preference

As a public-sector institution, GenomeDK must run a competitive tender every time it buys hardware and cannot simply repurchase from the same vendor. BeeGFS's hardware independence makes that model workable: each tender can go to whichever compliant hardware is cheapest.

Results

- A 3.5-person team operates a 33 PB, ~1300-user, multi-thousand-job HPC storage environment without the fragility of their previous enterprise platform.
- Public-tender hardware procurement is fully viable because BeeGFS is hardware-agnostic - driving down costs without compromising the storage platform.
- Built-in observability tools let a small operations team pinpoint which of thousands of concurrent jobs is degrading performance for everyone else.
- One architecture supports both active analysis and 15-year data retention, expanded incrementally in line with a fixed annual public budget.
- More than a decade of continuous BeeGFS operation - reliability the team sums up simply: "it pretty much just works."

Looking Ahead

GenomeDK expects to keep expanding its BeeGFS footprint, adding roughly 10 additional petabytes over the next two to three years as part of its ongoing incremental model, alongside a broader wish for richer statistics and diagnostic tooling to speed up identifying which job or user is behind a given performance issue.

“We don't have any plans to search for alternatives; the plan is to stay with BeeGFS.”

Dan Søndergaard, HPC senior consultant, GenomeDK

Trippstadter Str. 113
67663 Kaiserslautern, Germany
Phone: +49 631 277 57 63 00
Email: info@thinkparq.com
www.beegfs.io