

# Genomics England Improved Scale and Performance for On-Premises Cluster

## Challenges

- Scale capacity in a single namespace to store 140 petabytes of data
- Improve storage performance to accelerate innovation
- No strategy to protect valuable national data from a major disaster
- Tight budget constraints

## Benefits

- No limit on capacity scaling
- 10X improvement in performance
- 75% reduction in storage cost per genome
- Full disaster recovery strategy in place
- Integrated with public cloud for compute elasticity

NeuralMesh™ by WEKA® allows Genomics England (GEL) to scale to extreme performance and capacity in their mandate to sequence 5 million genomes by 2023. GEL, headquartered in London, is owned by the UK Department of Health and Social Care to run the 5 million genomes project, announced in 2018, which aims to sequence 5 million genomes from National Health Service (NHS) patients with rare diseases.

The organization acquires sequenced DNA samples from NHS patients and allows a team of over 3,000 researchers to use this data for medical research. Having completed the sequencing of the first 100,000-genome project, GEL has already acquired 21 petabytes of genome data and it is projected to grow to over 140 petabytes by 2023. The research conducted requires access to the entire dataset and must allow researchers to query the data in a highly randomized fashion. Therefore, all data must be stored in a single storage system.

# The Challenge:

## Poor performance at large capacity scale

In 2018, Genomics England approved a new storage platform to support the 5 million-genome projected growth through 2023. GEL had implemented a scale-out NAS solution from a leading vendor to support the 100,000- genome project, however it had already hit its limit on storage node scaling, and performance suffered when the system was near capacity. In addition, the existing solution had no viable disaster recovery strategy, as backing up all 21 petabytes of storage is infeasible. This left key national data in a vulnerable position should a major disaster occur.

The GEL infrastructure team determined they needed a new storage strategy to support the anticipated growth through 2023.

“

We needed something that's much more scalable than existing NAS solutions – an infrastructure that could grow to hundreds of petabytes. Our existing solution couldn't provide that scale and wasn't performing as well in these magnitudes – that's what drove us to WEKA.

*David Ardley,  
Director of Infrastructure  
Transformation, Genomics  
England*

The new solution had to meet several key criteria:

- Scale in line with the anticipated growth to 140 petabytes in a single storage system
- Meet demanding performance requirements of the bioinformatics pipeline and core research
- Provide a disaster recovery solution
- Be as easy to manage while delivering required enterprise features
- Fit within the allotted budget

# The Solution:

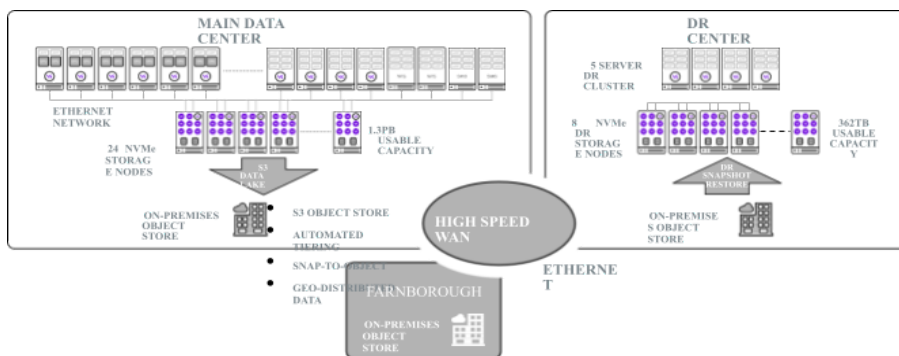
## Neuralmesh by WEKA<sup>®</sup> on industry standard server infrastructure

GEL evaluated all the major contending solutions through a rigorous RFP process. They rejected parallel file systems due to the complexity and lack of enterprise features. They rejected all-flash scale-out NAS because the economics were not viable for a project at the scale of GEL's requirements. Ultimately GEL chose WEKA because it was the only vendor who could deliver a solution that met all the requirements in a single architecture.

NeuralMesh is a fully parallel and distributed file system that has been designed from scratch to leverage high performance flash technology and cost-effective disk storage. Both data and metadata are distributed across the entire storage infrastructure to ensure massively parallel access to NVMe drives. Data is seamlessly tiered from flash to disk with NeuralMesh's internal tiering mechanism, allowing for best use of storage media and best economics.

NeuralMesh delivered a two-tier architecture that takes commodity flash and disk-based technologies and presents it as a single hybrid storage solution. The primary tier consists of 1.3 petabytes of high performing NVMe-based flash storage, which supports the working datasets. The secondary tier consists of 40 petabytes of object storage to provide a long-term data lake and repository. NeuralMesh presents the entire 41 petabytes as a single namespace.

Each of the tiers can scale independently: Should GEL require more performance on the primary tier, it can do so independently of the data lake. The system takes advantage of the geo-distributed capability of the object store, and data is protected across three locations, 50 miles apart. If there is a major disaster in the primary location, NeuralMesh's unique Snap-to-Object feature allows the system to be restarted in a second location, ensuring continued access to the data. The geo-distribution and object store resiliency provide significantly higher data protection, as full backup is impractical at this scale.



Genomics England Infrastructure with Disaster Recovery

# Benefits and return on investment

GEL realized several benefits and tremendous return on investment by choosing NeuralMesh:

- GEL enjoyed a 10X+ increase in performance compared to their legacy NFS-based NAS. The current system is capable of delivering over 135 Gbps from the NVMe tier and performance will continue to scale as the cluster grows.
- The storage cost per genome has dropped from £52 to £13, a 75% drop in storage cost, and is anticipated to reach £2 by 2023, a 96% cost reduction. This has been accomplished due to the ability to integrate low cost disk-based object storage into the solution.
- GEL can survive a major disaster at the primary site and still maintain access to the data. The object tier is geo-distributed across three sites, each 50 miles apart. Should the primary site fail, a small disaster recovery cluster is available on a second site and can re-hydrate the file system from the most recent snapshot.
- Critical datasets are fully encrypted from the high-performance compute cluster all the way to the permanent data store, with integration to a key management system. In addition, the system is protected from rogue security threats through a robust authentication mechanism.

Finally, new opportunities have opened up for GEL with NeuralMesh. As more research comes online, GEL would like to integrate public cloud for compute elasticity. NeuralMesh will allow GEL to burst to the cloud and leverage on-demand compute resources.

## Huge reduction in storage cost per genome

