

CASE STUDY

EIDF enables world-leading data-driven innovation with storage solution from Spectra Logic

We are creating what we believe is the only facility of its kind in Europe focused on delivering data-driven regional growth. Spectra's TFinity is uniquely positioned to provide the scalable infrastructure for our long-term archive and storage services offered at the EIDF. It builds on a great long-term relationship we've had with Spectra.

Mark Parsons, Director of EPCC, University of Edinburgh

Spectra TFinity Tape Library



About EIDF

The Edinburgh International Data Facility is the set of data and compute services, organized as a Private Cloud that underpins the Data-Driven Innovation Programme at the University of Edinburgh. The programme is helping organizations tackle challenges for industry and society by doing data right and supporting Edinburgh in its ambition to become the data capital of Europe. The EIDF provides a rich environment of Data Science, AI computing and long-term data management services for its many partners from the public, private and academic sectors.

The Challenge

EPCC is the UK's National Supercomputing Centre. The Edinburgh International Data Facility (EIDF) is located at EPCC's Advanced Computing Facility in a purpose-built high-resilience computer room. The EIDF is possibly Europe's first regional data innovation centre and is based at the University of Edinburgh in Scotland. EIDF brings together regional, national and international datasets to create new products, services and research. Services range from private cloud compute to data management to safe haven services and bespoke projects. EIDF computing services are designed for data science, including its GPU-enabled machine-learning compute platform for driving artificial intelligence (AI) application development, focused on Europe's first Cerebras CS1 large-scale AI engine.

EIDF data management services enable users to archive and re-use data, offering long-term storage and preservation for the Data Driven Innovation (DDI) programme's data assets while growing EIDF's analytics-ready collection to make data available for researchers and innovators. EIDF also provides a large trusted research environment that can be tailored to meet the information governance needs of data controllers. The EDIF provides a range of secure environments for data analysis and archive services. Finally, EIDF works



Aerial shot of the development of the Edinburgh International Data Facility and the building that will host it at EPCC at the University of Edinburgh.

CASE STUDY: Edinburgh International Data Facility (EIDF)

with partners from the DDI programme to co-create and tailor services using a building block approach to data hosting, high-performance AI analytics, modeling, and beyond.

To support their full range of research services, EIDF required a data storage solution that enabled them to be as agile and responsive as possible to the needs of their many stakeholders across the DDI programme, which involves collaboration between industry, the public sector and academia.

The Solution

To set the foundation for their long-term storage strategy, EIDF implemented a joint solution from HPE and Spectra Logic. EIDF deployed two four-frame Spectra TFinity ExaScale Tape Libraries with LTO-8 tape drives and media. The TFinity systems are uniquely designed with robotic pickers called High Performance Transporters each tape library contains two to support dual operations. Instead of individual tape slots, the library has shelves of TeraPack containers - ten tape cartridges bound together in a single tray. The TeraPacks allow for the library's higher tape storage density in a smaller footprint than would be possible with a traditional design. This compact architecture allows the robots to travel a linear path. loading and unloading multiple tape cartridges at a time for maximum efficiency. EIDF uses HPE Data Management Framework (HPE DMF) software to enable hierarchical, tiered storage management to the TFinity tape libraries, optimizing data accessibility.

To manage AI workloads and applications at scale, EIDF is exploring HPE's Ezmeral Container Platform running on HPE Apollo Systems. The HPE Ezmeral Container Platform provides native Kubernetes support and enables self-service artificial intelligence/machine learning (ML) applications for EIDF scientists, with flexible use of accelerators such as GPUs – allowing developers to standardize ML workflows. EIDF has deployed a mixture of HPE storage solutions including Lustre and CEPH for persistent data storage for high-performance and high-throughput advanced analytics.



EIDF is combining computing and data resources to create a facility that will allow users to leverage data to innovate throughout their organizations. The TFinity tape systems provides the scalability that this project demands as it develops over the next decade. **J**

Mark Parsons, Director of EPCC, University of Edinburgh



The combined storage solution enables EIDF to meet their most challenging longterm scalable storage requirements. EIDF has 20 petabytes of cold storage capacity with the new system which is being used for a multitude of EIDF projects including vital COVID-19 research at the University of Edinburgh.

Environment Snapshot

- Two four-frame Spectra[®] TFinity[®] ExaScale Tape Libraries
- LTO-8 tape drives and media
- BlueScale® standard encryption
- Cerebras Systems Al Supercomputer
- HPE DMF software
- HPE Apollo Systems
- HPE Superdome Flex Servers
- HPE Ezmeral Container Platform
- · 27 petabytes of mixed fast storage

Why Spectra?

- · Consultative solution design
- Scalable and adaptable to growth
- Integration with HPE
- Excellent customer and product support
- Superior reliability, scalability and accessibility

Solution Recap

Spectra TFinity ExaScale Tape Library – With unsurpassed storage density in the smallest footprint of any enterprise library on the market, the Spectra TFinity ExaScale offers industry-leading scalability with the speed necessary to meet requirements of the most demanding environments. In addition to LTO tape technology, the Spectra TFinity Exascale is also compatible with IBM's TS11X0 enterprise tape technology and Oracle's T10000x enterprise tape technology, enabling all three tape formats in the same library.

©2022 Spectra Logic. All trademarks and registered trademarks are properties of their respective owners.