

## Case Study

# Accelerating data analytics workloads for Deutsche Bank



## Deutsche Bank

The Finance Data Warehouse (FDW) group at Deutsche Bank is responsible for mission-critical execution of analytics workloads required for regulatory compliance and ensuring the bank's stability.

### The Challenge

FDW found itself challenged to meet regulation deadlines due to an outdated monolithic solution. Revamping and distributing the workload required an overhaul of the underlying infrastructure, with an emphasis on the storage layer. FDW infrastructure was based on a separate instance of SAS statistical analysis software for each compute node, which communicated independently with external storage. With several environments, ranging from development to production as depicted in Image 1 below, this configuration was expensive and inefficient resulting in high licensing costs. Furthermore, as the input data volume increased from 22 to 33 million transactions over the course of the system's life, daily runtime increased, which made meeting the regulation deadlines impossible.

The team performs numerous routine calculations and ad-hoc reports such as group-wide consolidation, credit risk values and stress tests. To meet regulatory compliance, these reports and tests, which consume high data volumes, often exceeding 1 TB per run and a great number of variables, must be prepared on a daily basis within a constrained time frame. Storage high performance and scalability are essential features.



**Image 1**, typical mission-critical financial systems are built with several configurations to ensure production quality

## The Solution

To improve and revolutionize the way their applications function, FDW has created a common layer for all nodes, through a hyper-converged state-of-the-art cloud-oriented architecture. Excelerero's NVMesh software-defined storage was laid over the hardware to create a protected storage layer pooling high speed NVMe drives instead of a legacy appliance. IBM Spectrum Scale was placed on top of Excelerero NVMesh to provide a shared file namespace enabling the distributed SAS layer to function optimally and allowing it to scale as workloads naturally grow over time. The new architecture is shown in Image 2.

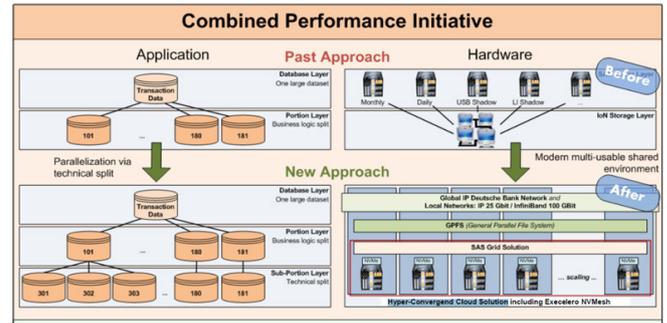


Image 2, moving to a cloud architecture reduces license costs, provides scalability and reduce turnaround times

## The Outcome



Low latency and high throughput allow finishing analytics workloads much faster Scalability, scaling the system is as easy as adding new nodes or new drives



Less hardware means less cost



**72%** savings in production run times, i.e. the application runs **3.5** times faster, which means achieving critical regulatory compliance



No proprietary storage appliances provide solution longevity



High utilization based on distributed storage and efficient CPU consumption



Initially, double the amount of workspace and file system space compared to the previous architecture with the assurance of scaling without further revamping