PRESIDIO® FEDERAL

Expanding a National Labs IT Environment with IBM Storage

Presidio Federal is proud to be a trusted advisor to many national labs, both small and large, that make up the U.S. Department of Energy (DOE) and to be actively engaged in the planning and design of their various technologies. Our dedicated DOE team's priority is providing the right technology for the complex DOE environment and pre and post sales service and support. Presidio Federal is optimized to deliver the industry leading technologies that are the product of IBM's continuous and successful, multi-billion-dollar annual investment in research and development (R&D).

The Presidio Federal DOE Account Management team has decades-long relationships having sold numerous IBM storage solutions throughout the DOE. In 2021, Presidio Federal placed over 25 petabytes (PBs) of usable capacity in three different Elastic Storage Server (ESS) clusters for DOE labs in the US.

The Presidio Federal DOE Account Management team has decades-long relationships having sold numerous IBM storage solutions throughout the DOE.

Presidio Government Solutions | 12100 Sunset Hills Rd #300A | Reston, VA 20190 | T: 571-789-0390
www.presidiofederal.com

These clusters are designed and implemented to support high performance computing (HPC) and artificial intelligence (AI), machine learning (ML), deep learning (DL) shared storage services to support large compute clusters with thousands of nodes and hundreds of engineers. Two of the clusters replaced previous generation ESS solutions that were recommended by the same account manager. While the customer was highly satisfied with the older ESS, they knew that the next generation ESS was the only storage platform to support current and future storage requirements.

This story of evolving, multi-generational ESS customers is common among IBM ESS customers who have large unstructured data processing and management requirements. They need and appreciate the ESS platform's value and ability to support their agency's mission on many levels: high performance at scale, high reliability at scale, and high flexibility to support a multitude of different workloads in a single name space or distributed global name space.

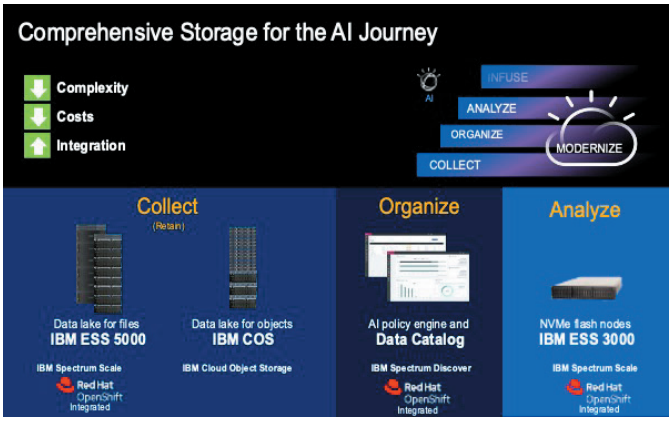

IBM ELASTIC STORAGE SERVER – HPC/BIG DATA/AI STORAGE PLATFORM OF CHOICE

IBM's Elastic Storage Server, with IBM Spectrum Scale (formerly known as GPFS) as the file system, provides non-disruptive modular and linear scalability of capacity and performance (starting as small as a 30TB of useable capacity and scale to hundreds of PBs). For example, a very large ESS cluster capable of supporting the most demanding HPC and AI workloads is the IBM ESS Alpine 250 PB file system supporting Oak Ridge National Lab's IBM Summit Compute Cluster, which is #2 on the HPC Top 500 list of worldwide Super Computing systems. The ESS makes it simple for HPC and AI/ML/DL teams to start small and grow with confidence. With performance, scalability, reliability, policy driven data management, with multi-protocol support and support for the new Container Storage Interface (CSI) open standard as the primary design points, the IBM ESS is purposefully designed and built to support modern cloud deployment models, Hadoop, HPC, as well as traditional systems/ solution architectures - on one platform.

Organizations can now truly benefit from the promises of shared storage service for all mission critical big data, HPC, AI/DL/ML and Hadoop workloads.

NO MORE DATA MOVEMENT

Access to the same data set can be done simultaneously, via multiple protocols from NFS, SMB, S3, Rest APIs, CSI, and HDFS. Enterprise class and proven data management capabilities allow for sophisticated policy driven replication, business continuance, data migration, and multi-site, global name space, and distributed data architectures.

Policy driven tiering from flash to tape to cloud has been supported for years. Encryption is supported with IBM Security Guardium Key Lifecycle Manager. Data reduction is accomplished with onboard compression with no performance impact. All processing, mid-term storage and long-term archiving can be supported with in a single ESS/Spectrum Scale global namespace providing universal access to data and performance. All open-source AI/ML/DL frameworks can run against these systems, as well as big data analytics allowing you to enhance the data science productivity.

PRESIDIO FEDERAL

IBM's ESS storage provides the perfect combination of performance, scalability, reliability, data management automation with integration with the rest of the IBM hardware and storage software products. The most rigorous and complex data processing, security, availability, scalability/flexibility requirements can be easily satisfied. The largest and most sophisticated supercomputing solutions in the world run IBM ESS storage ensuring their massive systems continue to run at peak performance!

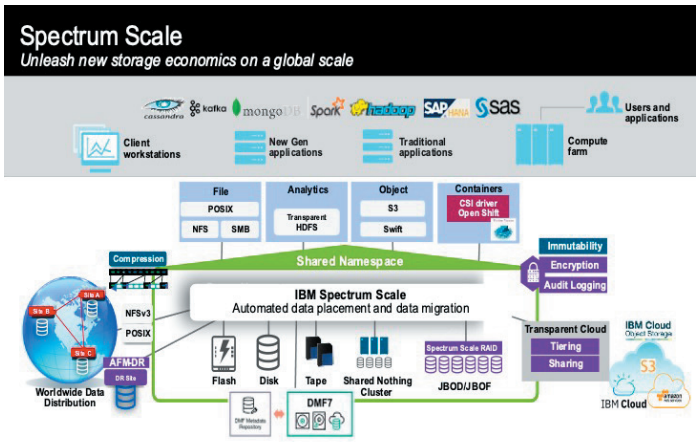

THE CHALLENGE

The need for all DOE national labs to enhance and continue to improve their storage platforms, comes from the expansion of the workloads and the increased need for virtual capabilities.

Systems must be economical and provide robust security. At the same time, they must provide shared storage that is designed to fulfill large, complex processing requirements that can be managed easily and holistically.

THE SOLUTION

With IBM's Spectrum Scale ESS integrated storage appliance, they were looking for a product that:

- Allows consolidation of numerous siloed storage systems into one storage platform simplifying the management of huge amounts of data
- Supports any performance and scale requitements
- Provides high RAS (reliability, availability, serviceability, data resiliency) levels – five 9's (99.999%)
- Minimizes software licensing through the elimination of silo storage systems
- Minimizes training requirements by using one common management interface
- Provides end-to-end optimization of the solution ensured by IBM for risk mitigation

IBM is the top storage recommendation as "it's tried and true" for DOE data at scale missions, providing a set of solutions that can address any storage requirements no matter the performance, scale, cost efficiency, density, flexibility, security, and availability requirements. No other storage vendor has the storage (hardware and software) portfolio, strategic vision, technologies, R&D spend, economy of scale, and the pedigree, which customer directly benefit, then IBM.

For current DOE ESS customers, IBM's ESS and Spectrum Scale provide the "easy button" to solve the most rigorous unstructured data storage requirements with high capacity, disk-based models, and high performance/ ultra-low latency flash-based models, which can be mixed and match with each other or within the same model to provide market-leading performance, flexibility, density, and scalability. The ESS is specifically designed to effectively support modern HPC and AI/ML/DL/Big Data compute platforms and supporting technologies, at scale. IBM Spectrum Scale can be deployed with the ESS, in the cloud, or on third party hardware and still work together in a global name space, with visibility to the data across all deployment modes.

For current DOE ESS customers, IBM's ESS and Spectrum Scale provide the "easy button" to solve the most rigorous unstructured data storage requirements.

PRESIDIO® FEDERAL

IBM's ESS and Spectrum Scale benefits include:

Optimal **building blocks** for high-performance, scalable, reliable enterprise Spectrum Scale storage. This includes the ability to scale up or scale out an agencies storage to meet the performance demand of their workloads. The IBM Spectrum Scale integration simplifies storage administration for day-to-day operations with robust graphical user interface that supports all administrative command-line interface (CLI) functions and Ansible automation scripting.

High Availability and Data Resiliency with highly flexible synchronous or asynchronous replication, snapshots, and numerous Business Continuity, Disaster Recovery strategies supported. ESS implements the unique advantages of IBM Spectrum Scale RAID, which is an erasure coding RAID strategy, which is the core of the Spectrum Scale file system, that spreads the data and parity data across all the drives in an ESS solution (which is highly flexible, not all or nothing).

Performance at hyper scale, fully automatic data management capabilities, and flexible integrated storage solution.

Spectrum Scale RAID, a sophisticated erasure coding technic to spread all data and parity across all drives in a given ESS solution, **greatly reduces the impact and risks of pact of large capacity drive failures**, providing significantly faster rebuild times (important for large disk drives), while providing consistent high performance to support a high number of aggregated and diverse workloads.

Single repository of data with unified file and object support. There is just one global namespace to expose to all users. Spectrum Scale provides multi-protocol support using protocol nodes, such as Network File Systems 4.1, Server Message Block, S3, HDFS, and a direct Spectrum Scale Native Driver for highest performance.

In millions of hours of production usage, ESS **reliability** has demonstrated five 9's (99.999%) of availability.

Presidio Government Solutions | 12100 Sunset Hills Rd #300A | Reston, VA 20190 | T: 571-789-0390
www.presidiofederal.com

PRESIDIO® FEDERAL

BENEFITS AND A LOOK TO THE FUTURE

Overall, data architecture is the biggest difference when comparing the before-and-after of implementing IBM's Spectrum Scale. ESS allows users to easily consolidate data with improved performance, RAS, and economy by keeping all the data in a single ESS global name space. This reduces network cycles, administration effort, and improves capacity utilization. IBM Spectrum Scale and ESS's data architecture allows all processing, mid-term storage, and long-term data storage in one logical Spectrum Scale file system.

The value of Spectrum Scale and ESS is simple: High-performance parallel data access with enterprise data services, supporting simultaneous access via multiple protocols to the same data from a single cluster that is built with simple scalable building blocks optimized for maximum capacity, density, throughput, low latency, and cost optimization. One storage platform for all your data and workloads.

This makes IBM's ESS and Spectrum Scale second to none for any AI/ML/DL, big data analytics, and any HPC workloads. Presidio Federal continues to work with many DOE labs with planning new ESS/Spectrum Scale clusters and showing how investing further into IBM ESS storage will benefit their mission.

ABOUT PRESIDIO FEDERAL

Presidio Government Solutions LLC, branded publicly as Presidio Federal, is a purpose-built and mission-driven IT services and solutions provider dedicated to serving the federal government. Presidio Government Solutions leverages its wealth of experience and deep relationships across its partner ecosystem, creating an environment of active collaboration and real-time responsiveness.

The company develops and delivers the most advanced technologies through expert knowledge centers in automation, augmentation, cloud, cybersecurity, digital infrastructure, and collaboration.

Presidio Government Solutions is a wholly owned subsidiary of Presidio Networked Solutions.

Presidio Government Solutions | 12100 Sunset Hills Rd #300A | Reston, VA 20190 | T: 571-789-0390
www.presidiofederal.com