

ExERA

Exabyte - Era - Archive

Best-of-Breed

ExERA - Overview

ExERA is an on-premise, purpose-built, engineered data archive solution specifically developed for HPC customers. The design principal centers around an architecture that delivers frictionless scalability in 4 dimensions: throughput up to 100 GB/sec, 100s of thousands of IOPS, high parallelism including a parallel file system, and beyond Exabyte capacity. ExERA is a robust, scalable architecture that provides flexibility and future-proofing as HPC data management and scalability requirements change dramatically over the next few years. Migrating to ExERA today may dramatically collapse workspace island infrastructure and position our customers to meet the challenges of rapidly growing archive capacity requirements at the Exabyte level.

ExERA is purpose-built and engineered solution that delivers:

- Scalable IO throughput to support current and future requirements.
- Automatic movement of data to the correct archive storage target, based on policies and budgetary considerations.
- Operationally-resilient environment that has high availability and fault tolerance.
- High volume, bi-directional data movement from and to the HPC primary scratch storage and secondary processing to the workspace and archive layers to support the emerging AI, ML, DL, Big Data, Data Analytics, Data Mining and Data Lake requirements.
- Future proofing by leveraging commercial off-the-shelf (COTS) with selected best-of-breed components and products. No custom hardware or software.

Storage Expertise

- Delivering years of HPC experience, DoE & DoD proven solutions with expert level Storage Solution Architects.

Industry Leading Solution

- Offering vendor-neutral, but best-of-breed components to manage billions of Files & Objects and Exabytes of data, while remaining open to emergent storage technologies.

Lowest Risk & Lowest Operational Costs

- ExERA, an on-premise solution, uses a Purchase, Lease or Cloud-based pricing model, designed to beat the large Cloud vendor's costs, even Glacier.

Next Leap in Archive Systems

- 5-10x traditional throughput and IOPS.
- 5-10x the capacity of traditional archive systems.

Alliance won the IBM Beacon award in 2021 for the ExERA solution.



ExERA



Throughput

Workspace design up to 100 GB/sec of ingest and recall.



Capacity

Non-disruptive, scalable performance to Billions of files, objects and Exa-scale capacity.



Secure

Designed for very long retention periods of 10s of Years or more.



Accessibility

Standard access protocols like NFS, SMB, CIFS, S3, Swift and FTP.



Reliable

Fault tolerant and high availability, no single point of failure architecture.



Affordability

Tape is still the lowest cost \$/PB even today and has a great technology roadmap.

Enterprise COTS Components

- ExERA uses proven best-of-breed COTS components from vendors in the HPC marketplace.

Manageability

- Tools to seamlessly manage performance storage policies.

Integrated

- Ability to assimilate into current and future HPC network topologies.

Flexible Delivery Model

- Purchase, Lease or On-Premise as a Service.

Purpose Built

- Custom architected from customer driven- requirements.

Partners

- Our goal is to incorporate our HPC partner's best-of-breed products into our solution as necessary to meet all customer requirements.

ExERA - Overview

One of the ExERA architecture design goals is the ability to create an on-premise cloud- like archive sharing for long term retention of files and objects. ExERA allows customers to establish both private, hybrid and public clouds archives to meet different operational requirements.

ExERA is about providing choices and the flexibility to meet individual customer requirements. However, these options and choices do not require performance tradeoffs. ExERA in fact, offers tremendous performance to meet the Throughput, IOPS and Scalability challenges through parallelism. For example, our workspace can support 10s of PBs and can easily support a 100 GB/sec ingest and recall processing and can scale from small to larger. Our use of flash for all the metadata provides msec access to internal file and object structures. Our use of flash for small files, like 100 TB or more of data space, provides the ability to support 100s of billions of files and objects with fantastic IOPS performance. The highest performant systems use all flash. Less performant requirements use a more traditional disk cache architecture with Hybrid flash and HDD arrays.

Parallelism:

- We design for a scale-up or scale-out set of x86 data movers from DellEMC, HPE or SuperMicro. ExERA is the solution that offers a way to save substantial dollars by collapsing or shrinking away at least one or more, secondary and post-processing workspace infrastructure(s). By leveraging our dual-use PFS Workspace layer our customers might be able to collapse some of the secondary processing IO storage islands currently installed. Marketplace examples are Hadoop, commodity storage, campaign storage, emerging AI, ML, DL, Big Data, Data Analytics, Data Mining, Data Lake and others.

Throughput and IOPs:

- ExERA's very high performant Workspace layer is supported today by the market leading parallel file system, IBM Spectrum Scale (GPFS) based on ESS arrays. For less demanding archives we use a high performant disk cache solution using one of many vendors' mid-range hybrid disk arrays. The benefit of the ExERA architecture is for heavy workload and processing requirements, this design can be scaled-out, by simply adding additional building blocks for both the throughput and disk capacity.

Capacity:

- ExERA offers four main archive layers, Disk, Object, Tape and Cloud. ExERA's high capacity Disk Archive Layer is supported by many enterprise-class Object Storage systems from vendors like IBM, DellEMC, Scality, Quantum, Netapp, Pure, others. The benefit of Object Storage is that additional throughput and capacity can be added within minutes to meet growing workload requirements. By leveraging the Disk Archive Layer, customers may be able to collapse some of the bulk storage islands like Commodity and Campaign. ExERA's high capacity Tape Archive Layer is supported by all industry-leading tape libraries from IBM, Spectra, Quantum and Oracle. The benefit of tape is that it is still the lowest cost per PB in the industry and allows for up to an Exabyte or more of capacity with current technologies. Solution costs as low as \$0.0001/GB/MO can be achieved.

The ExERA Solution from Alliance is significantly more than just a product vendor selling individual components like the archive software or just the tape library with tape drives. Our complete Archive-designed ecosystem includes networking, data movers, archive software, either an internal workspace or disk cache layer, possibly an object storage system, supporting POSIX, NFS, SMB/CIFS, S3 and tape library and tape drives. Marketplace examples might be IBM HPSS with an IBM library, aging Oracle HSM with aging Oracle library or just aging Oracle library, Quantum StorNext with a Quantum library, HPE/SGI DMF with a Spectra Library or Versity with any library.

Our proposed solutions can range from a single component replacement, for example library replacement or just software replacement, a partial solution replacement like software with data movers, or a complete takeout design replacement of everything. Some takeout designs involve no immediate media migration since the archive software selected can ingest the legacy metadata directly and can read the legacy media format.

The ExERA solution meets and exceeds the expected operational and performance requirements of many HPC customers, even those at Exascale. In fact, ExERA was designed based on direct input from HPC customers. ExERA with its very high performance workspace layer can handle tens of PBs with over 100 GB/sec for ingest and recall throughput. The Disk Archive layer has a target range of up to 100 GB/sec and hundreds of PB of object storage. The Tape Archive layer supports an Exabyte (EB) or more of capacity and up to 100 GB/sec of throughput. All layers have the ability to start small and grow very large over time.

ExERA provides individual OEM's for support, warranty or break/fix activities. The Archive Solution is typically custom-built on-site and tested, or we can pre-configure it in Alliance's secure facility near Baltimore, MD if needed. This is the same Alliance organization that provides world-class support to many organizations including the Department of Defense, the Department of Energy and others.

ExERA

Best-of-Breed

This layout and this partial partner list illustrates some of our HPC partners' portfolio components that can be optimized within the ExERA solution: current products can include Spectrum Scale (GPFS) or Lustre parallel FS software, IBM ESS, DDN EXAScaler, HPE ClusterStor parallel FS arrays, IBM, DellEMC, HPE, Netapp, Pure, others all flash and hybrid arrays, HPSS, Spectrum Archive, Stornext, DMF, Versity, Atempo archive software, DellEMC, HPE and Supermicro servers, IBM COS, DellEMC ECS, Scality, Quantum ActiveScale, Netapp StorageGRID, Pure, others object storage systems, IBM TS4500, Spectra TFinity, Quantum i6000, Oracle enterprise tape libraries, TS1160 & LTO9 tape drives and media, Brocade or Cisco FC switches, 100Gb/200Gb Infiniband and 25GbE/40Gb/50GbE/100GbE ethernet networking. ExERA delivers a fully configured Archive System with Professional Services to deploy layers for networking, servers, flash for metadata, disk arrays, object systems, libraries and tape drives. Of course external clouds. Growth uses pre-defined Building Blocks that may be non-disruptively integrated into each of the ExERA Layers for additional performance and/or capacity.

