elasticsearch + VAST

ELASTICSEARCH ON VAST

Elasticsearch is the foundation for the most demanding streaming-data and log-analytics platforms found in the wild. Its popularity has given rise to a large ecosystem of clients tools, use cases and visualization suites that have made it indispensable in the enterprise. Powerful analytic software requires powerful storage: Elastic deployments will bottleneck as data scales, forcing organizations to separate data into hot and cold tiers in order to manage costs. This compromises their ability to use all of their data, all the time, to generate value. VAST's goal of realizing all-flash performance with spinning media TCO and enterprise-grade reliability makes a tierless, performant and easy-to-manage Elastic infrastructure possible. No more compromises. Here's why.

VAST ARCHITECTURE

VAST's disaggregated shared-everything architecture (DASE) separates the stateless front-end servers, from the storage enclosures where state is stored. This allows VAST to scale compute and storage independently to precisely align with your workload with zero-downtime. In a standard Elastic configuration, each cluster node will mount a VAST NFS share unique to that node ID. This supplies a discrete slice of storage to each Elastic process, acting as a drop-in replacement for direct-attach storage in a standard architecture.

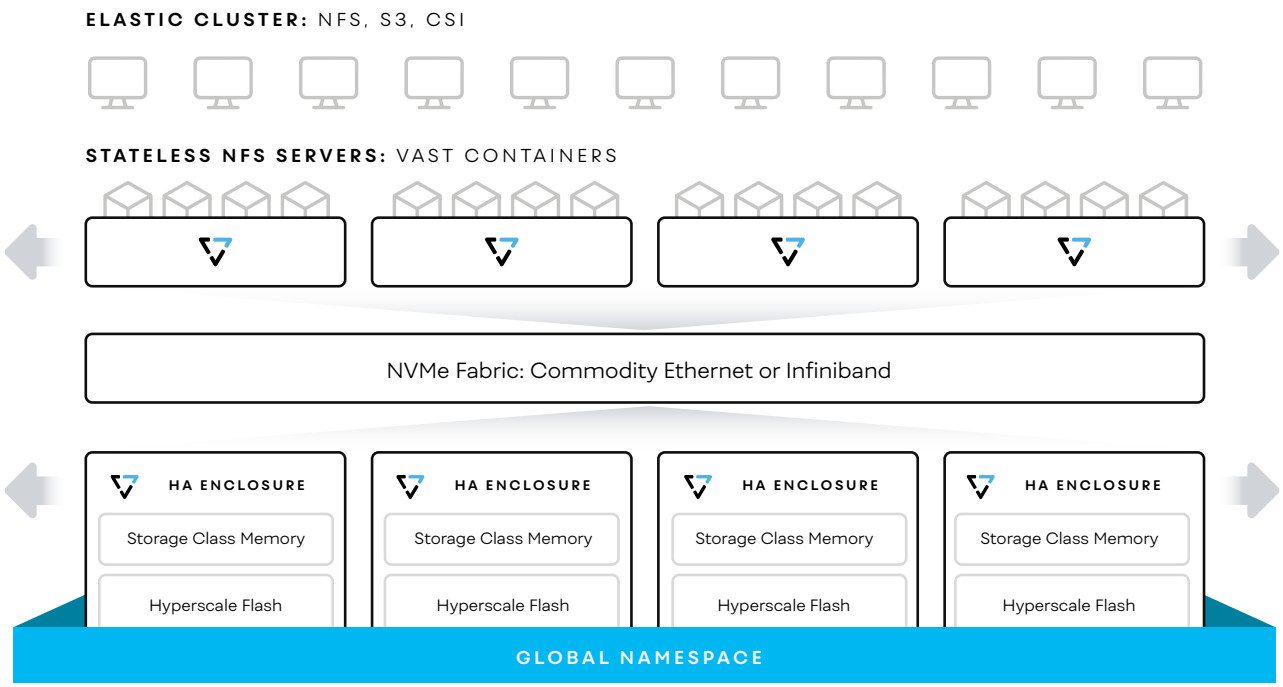

This works because VAST's high-performance NFS service supplies near DAS-like performance with the flexibility of a network share. Performance can be further increased on kernels with nConnect capability (available in any recent Linux distribution) supplying tens of gigabytes per second of throughput to a single Elastic node or ECK instance. Sites that want to take it a step further and get HPC-like performance can install VAST's RDMA+multipath NFS client.

Regardless of the NFS client decision, VAST supplies fast, resilient storage to your Elastic cluster.

TIERLESS ELASTIC ARCHITECTURE

A major factor in making a tierless flash-based architecture possible is VAST's advanced deduplication and compression techniques, resulting in a data reduction ratio that averages roughly 3:1 across our install base. Reduction that eliminates many of the cost barriers to creating an all-flash architecture: Now you have high-performance access to all of your Elastic data, all of the time, without compromising on retention periods.

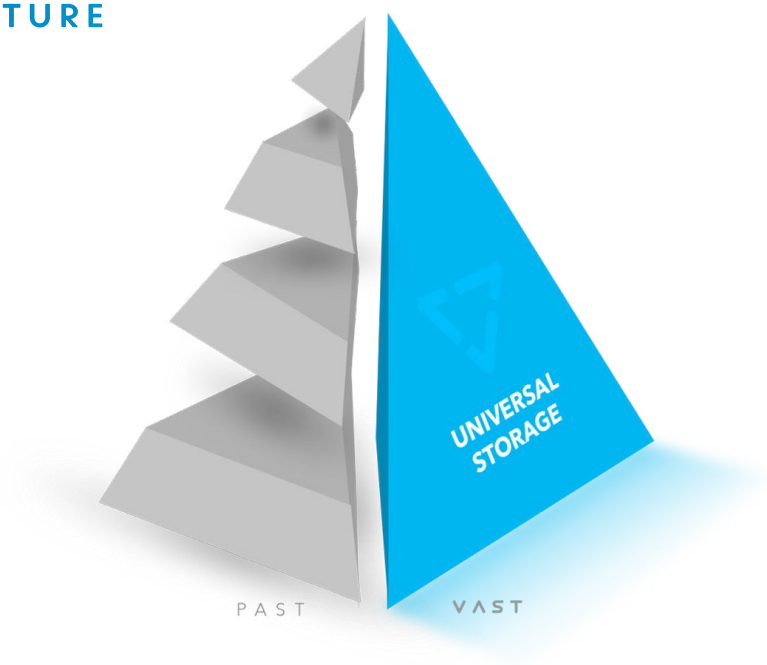

STORAGE EFFICIENCY AND COMPUTE SCALING

Elastic can use multiple replicas of data shards in the cluster to increase search parallelism. With traditional DAS, NAS or block storage systems, scaling replicas for performance can dramatically increase your storage footprint. VAST's high-performance in-line deduplication techniques effectively remove the penalties for storing additional copies. The Elastic cluster "sees" multiple replicas but the space consumed in VAST is effectively that of a single copy.

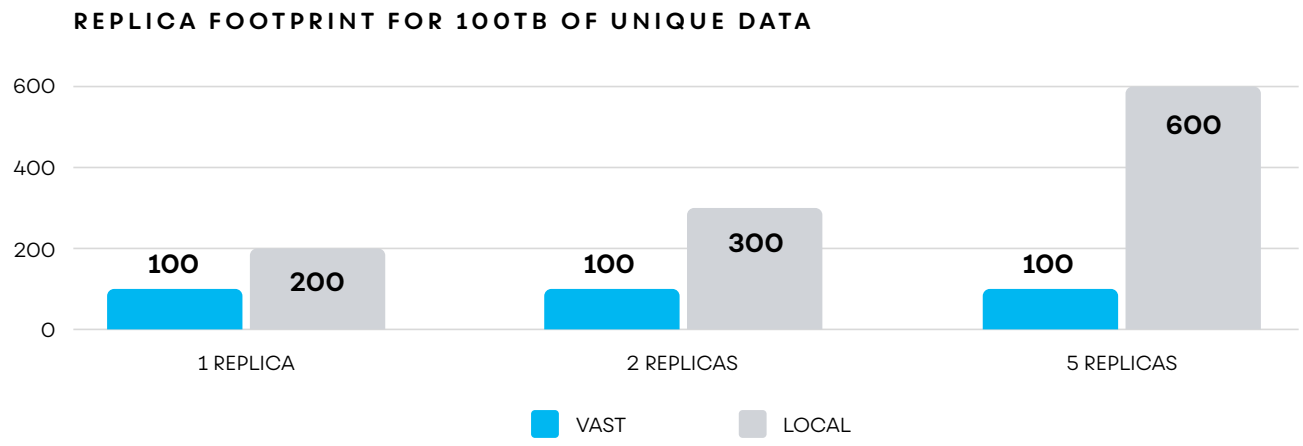

ELASTIC IN KUBERNETES (ECK)

VAST's CSI driver for volume management in Kubernetes can be combined with ECK to enable scalable Elastic deployments in cloud environments that are serviced by VAST. It opens up a range of options for your deployment to suit whatever availability or performance model you want.

FLEXIBILITY

VAST includes a peerless high-performance S3 capability alongside NFS and SMB. This means that when cold and frozen tiers are a structural requirement in your Elastic architecture, VAST can support all levels - hot, warm, cold, frozen - all on the same storage system.

All-flash storage, multi-protocol support, erasure coding, advanced deduplication, compression, CSI support and an enterprise-grade resilient storage architecture all come together at a cost-point to make VAST perfect for housing for Elastic data at petabyte scale.

Advanced compression and deduplication	Reduce storage footprint regardless of replica count, delivers unbeatable price/performance on flash.
High-performance NFS	Meet performance-level agreements and get real-time results
CSI Driver	Deploy Elastic in public/private clouds powered by VAST
DASE Storage Architecture	Scalable capacity/performance and 99.99999% availability