

Vizion.ai Elasticsearch Service and Scale

Last Updated 13 May 2019

Vizion.ai Elasticsearch Service (ESS) provides the easiest environment to run Elastic Stacks at scale. Built on a strong heritage of distributed data management at the petabyte level across millions of users, Vizion.ai ESS greatly simplifies the scale and management for Elasticsearch (ES). There are several design principles for Vizion.ai ESS: simplicity, scale, solutions, security, and savings.

- **Simplicity.** Vizion.ai ESS greatly simplifies Elastic Stack provisioning by automating the infrastructure process. While Elasticsearch is a distributed application, the underlying infrastructure is not. Aligning the right infrastructure resources to the right Elastic Stack workloads is a challenge. When users define their Elastic Stacks in Vizion.ai, they do not have to separately choose different nodes or other resources. All infrastructure selection and software deployment is automated. Automating this process allows customers to focus on data analytics instead of cluster management and maintenance.

With Vizion.ai ESS, Elastic Stacks are immediately created and provisioned without requiring users to specify infrastructure components.

- **Scale.** Vizion.ai ESS automatically scales underlying infrastructure for every ES workload for every Vizion.ai tenant. Vizion.ai extracts different infrastructure layers to constantly improve the performance and economics for every ES workload. When considering indexing data storage, such as web log data in Vizion.ai ESS, each ES data shard has hard limits on the number of documents it can contain. As new events are logged as new documents, the prevailing index template always creates new indices that follow existing mappings and queries that the new indexes require. While an ES within Vizion.ai is able to continuously access all of its data, the underlying infrastructure components have been abstracted. With all cluster resources being both vertically and horizontally scaled, typical cluster intervention is greatly reduced. Tasks that affect SLAs, such as cluster I/O rebalancing for a dynamic number of data or ingest nodes, are also greatly reduced.

Vizion.ai ESS does not change any aspect of how Elasticsearch indexes data. It reduces the trade-off between archiving and snapshotting data to off-line locations like separate HDFS or S3 buckets. This can potentially reduce the fidelity of time series data types with log roll-ups. In this example, we allow ES clusters to retain practically unlimited historical data while making all data immediately available for any query. Vizion.ai decouples compute, storage, and other infrastructure services so that workloads can be replatformed with the right resources without interruption.

- **Solutions.** The Vizion.ai ESS platform is API compatible with dozens of other data sources ranging from web log information from NGINX to streaming data via Kafka. Vizion.ai ESS has worked to preserve all API integration points to offer full compatibility with existing data ingestion for all Beats, Logstash, or any other JSON-based document ingestion. For visualization, each ES stack from Vizion.ai comes with a built-in associated Kibana presentation layer. Vizion.ai also provides the fully autoscale ES stack that can be accessed using other visualization tools, including Grafana or direct ES query.

- **Security.** Vizion.ai ESS builds-in security. As Vizion.ai ESS offers the simplest and most scalable ES solution, all security is naturally built-in from the start. From its security and privacy based DNA to its security-by-design and security-by-process implementation, Vizion.ai greatly reduces operational risks to its tenants. Every new tenant and every Elastic Stack is effectively isolated from each other. For more information about Vizion.ai's security, visit <https://go.panzura.com/rs/346-POI-893/images/Vizion-ai-EsS-SECURITY-PRIVACY.pdf>
- **Savings.** Vizion.ai ESS provides a significant cost savings in several ways. Like all public cloud and public SaaS services, Vizion.ai ESS takes advantage of economies of scale within its multi-tenant environment. The marginal cost to create new Elastic Stacks within Vizion.ai is negligible, and Vizion.ai passes this cost savings onward. Furthermore, Vizion.ai's scaling and infrastructure abstraction costs are lower from a CPU or uptime perspective. This is because infrastructure is always optimized to the workload without excess provisioning and waste. Vizion.ai saves administrator time by drastically reducing time spent on adjusting, reconfiguring or rebalancing clusters. Vizion.ai's infrastructure abstraction proactively finds and uses lower cost tiers of storage, compute, and other infrastructure when not needed. Vizion.ai ESS also proactively reduces any underused or spare compute, memory, or storage resources, passing along the savings to tenants without service interruption or degradation. Finally, Vizion.ai saves customers money by integrating backup and recovery features, which otherwise would need to be managed by the customer. Vizion.ai ESS was designed to save on both infrastructure operating costs as well as the most important resource - time.

In summary, Vizion.ai ESS and its design principles for scale, security, and savings allow tenants to analyze, search, find, and derive insight from their data faster. Vizion.ai invites and offers tenants to try Vizion.ai ESS and experience the simplicity and scale first hand. Visit <https://www.vizion.ai/> to get started today.