

Accelerating Enterprise Analytics with Vizion.ai Elasticsearch Service

Elasticsearch is increasingly becoming a key platform in next generation analytics architectures. As a resilient and distributed application, it enables rapid data ingestion, indexing, and visualization across a broad range of enterprise use cases. And as an open source platform it has a broad community of developers with a rich ecosystem of data plug-ins and visualization dashboards. The fastest growing segment for Elasticsearch is in the IT operations and analytics market, aka log analytics, which IDC estimates to be a \$2.2 Billion market growing at 20% annually. In many cases organizations are exploring Elasticsearch as an open alternative to more expensive proprietary logging solutions like Splunk and Sumo Logic.

While Elasticsearch is easy to download and deploy for small use cases of a few terabytes, scaling to very large enterprise-wide production clusters is costly in terms of configuration management, operational support, performance troubleshooting and upgrades. For illustration, below is E-Bay’s Elasticsearch management process.

While Elasticsearch is easy to download and deploy for small use cases of a few terabytes, scaling to very large enterprise-wide production clusters is costly in terms of configuration management, operational support, performance troubleshooting and upgrades. For illustration, below is E-Bay’s Elasticsearch management process.

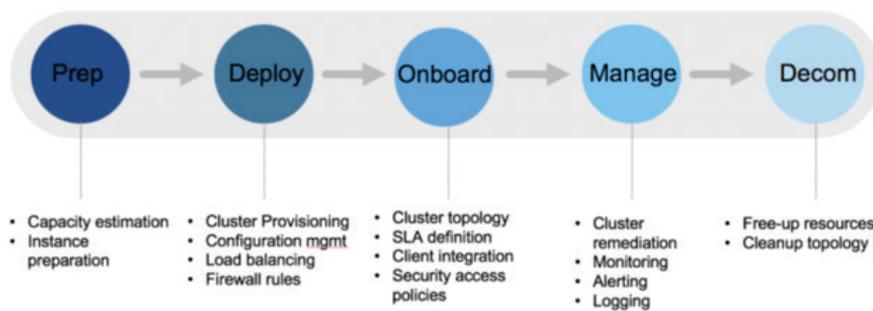


Fig 1. Traditional Elasticsearch Cluster Lifecycle Management process by E-Bay

Vizion.ai’s Elasticsearch Service was launched in the Spring of 2019 by Panzura to solve the challenge of running Elasticsearch at scale for the enterprise. Delivered as a fully managed service, Vizion.ai eliminates the management overhead and scaling limitations and provides a unique storage architecture that auto-scales the cluster to meet data ingestion and retention needs. Our proprietary Vizion Block Object Storage (VBOS) layer enables all indexed data to remain hot and searchable, but at a low-cost profile of object storage. This is increasingly important as ML and AI algorithms demand larger data sets for accurate training.

The Vizion.ai service delivers a new cost-effective consumption model for Elasticsearch across the enterprise while eliminating barriers of adoption to this new powerful data analytics tool. With a deep heritage in distributed data technology covered by 30 patents and over 7,000 deployments across 33 countries, Panzura has a proven track record of securing, protecting and managing enterprise datasets in the hundreds of petabytes.

Vizion.ai Elasticsearch Service (ESS) is implemented in a microservice architecture that abstracts the infrastructure layer and adds a block to object caching tier, thus simplifying the operation and scale for Elasticsearch. Vizion.ai is offered as a turnkey managed service deployed in a customer’s data center on commodity hardware for those seeking private cloud deployments. A multi-tenant public cloud version is also available at www.vizion.ai. Thus Vizion.ai can be consumed in a private, hybrid or multi-cloud manner.

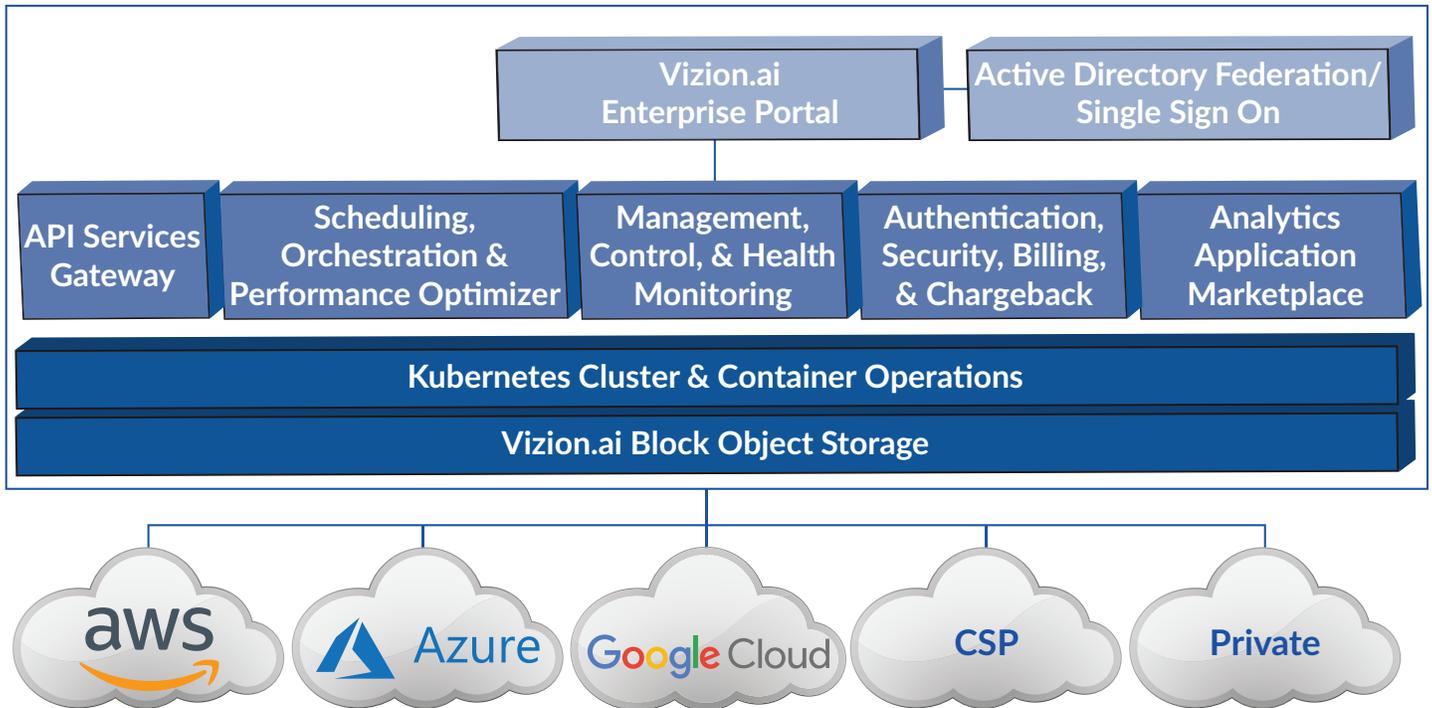


Fig 2. High Level Overview of Vizion.ai Architecture

There are four main use cases for Vizion.ai: security information and event management, application performance monitoring, IT operations & incident management and IoT/digital transformation. With hundreds of data connectors in the open-source Elasticsearch community, Vizion.ai offers fast time-to-value with out-of-the-box dashboards immediately available.

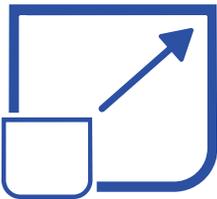


Fig 3. Operational Dashboards Bundled with Vizion.ai Elasticsearch Service

123

Simplicity

Traditionally as noted in Figure 1, when deploying an on-premise Elasticsearch cluster for log analytics it's necessary for IT admins to create server clusters, master nodes, indexing nodes and dedicated storage pools for the aggregated log size. The Vizion.ai platform eliminates this process. In just one click with Vizion.ai a developer can create a production-ready virtual cluster. This abstraction enables IT administrators and data scientists to focus on gaining value from the data rather than manage and maintain infrastructure.



Scale

Vizion.ai ESS allows scaling without limits. With the VBOS technology all storage pools scale elastically to accommodate the data set sizes driven by the application. This eliminates stranded storage or insufficient storage scenarios that can cripple Elasticsearch performance. Because the VBOS storage is a hybrid flash/object architecture, performance for hot data remains very high while maintaining a very low-cost storage profile and keeping all data online for analytics.

Vizion.ai's monitoring and automation engine varies the compute, storage, memory and network resources allocated across the cluster based on demand. This enables automated cluster rebalancing and shard migration based on capacity utilization metrics. As data ingestion fluctuates, Vizion.ai will ensure optimal performance that meets application needs in real time.



Savings

The savings offered in Vizion.ai comes from years of development in block/object caching technology enabling low-cost disk storage to perform like flash drives. By affordably keeping larger data set sizes, machine learning and other analytics tools can build better models resulting in higher accuracy at lower costs. No more having to move data to cold tiers where it is unsearchable. The net result is savings to the bottom line with improved IT operational insight and efficiency. Vizion.ai also saves operational cost by eliminating personnel time spent on adjusting, reconfiguring, rebalancing and upgrading clusters.

Vizion.ai is offered on a pay as you go, usage-based model. The pricing model considers data indexed, data stored, and queries – not the underlying infrastructure or node counts. Vizion.ai ESS saves organizations up to 50% compared to building it yourself, and typically 60-70% savings over proprietary solutions such as Splunk and Sumo Logic.



Security

Vizion.ai ESS is a microservice based architecture that operates on the basis of least privilege and reduces the attack surface area by removing any access to either operating system or runtime level components. Drawing from a full multi-tenant architecture, all tenant data remains fully isolated. All data at rest and in flight remains encrypted and isolated from other Vizion.ai tenants. By integrating with corporate Active Directory and single sign-on solutions, access to Vizion.ai can be tightly controlled with existing corporate identity and authentication mechanisms.

Open Architecture

Vizion.ai takes an open-source approach to Elasticsearch using the Open Distribution for Elasticsearch, An Apache 2.0-licensed distribution enhanced with enterprise security, alerting, and more. Backed by Amazon Web Services with supporters including Expedia and Netflix, the Open Distribution for Elasticsearch maintains a 100% Apache 2.0 open-source approach ensuring long-term application portability without licensing restrictions. With this open architecture there are hundreds of data collectors for sending logs into Vizion.ai including Filebeat, Auditbeat, Metricbeat, Packetbeat and Logstash that can send practically any log record or data stream from an IOT device, server, edge device, switch or operating system into Vizion.ai ESS.

Public, Private and Hybrid Cloud Deployments

Vizion.ai is offered both as a multi-tenant public cloud service and as a dedicated, private managed service that can be deployed on customer premises. For deployment on customer premises, there are multiple options available including installation into an existing Kubernetes environment or remote installation onto bare-metal servers. In all scenarios Vizion.ai provides full operational management, support, maintenance and upgrades.

Pricing and Contact

Vizion.ai provides transparent public cloud pricing that can be found at www.vizion.ai. For a typical log ingestion customer Vizion.ai provides a 60-70% savings over competitors such as Sumo Logic and Splunk. For private cloud customers savings can be greater with larger indexing workloads. For more information, questions or to schedule a demonstration, please contact us at info@vizion.ai.