

Accelerate Intelligent Search Performance with the Most Widely Deployed Vector Database

Elasticsearch delivers lightning-fast, accurate, AI-powered search across every data source

Organizations struggle with slow, inaccurate search that frustrates users and limits AI adoption. Modernizing search platforms helps teams erase technical debt and deliver precise, real-time results.

Power enterprise-grade AI with a unified search engine


Elasticsearch combines hybrid search, vector search, and semantic relevance in one platform. Integrated natively with Azure OpenAI, Azure Kubernetes Service, and Azure Blob Storage, the solution ignites search and accelerates Gen AI and RAG application development.

Unlock enterprise-ready speed, accuracy, and AI-driven relevance



Boost search speed at scale

Deliver sub-second search performance across massive datasets.



Increase precision and accuracy

Return precise, context-aware answers with clear citations.



Unify hybrid search methods

Blend keyword, vector, and semantic search in one engine.



Integrate AI and LLMs natively

Use inference with popular LLMs, built-in models and Azure AI services.

Scale AI-Powered Search with Flexible, Cloud-Ready Deployment

Elasticsearch provides a single retrieval layer for Gen AI, combining embeddings, filters, and semantic ranking with built-in vector search. On Azure, customers gain dynamic scaling, cross-cluster replication, and secure integration with Teams, M365, and Azure services.

Legora transforms document search with Elasticsearch intelligence

Confronting slow, fragmented search across complex legal documents

Legal technology provider Legora needed to modernize how its users accessed legal and regulatory content. Challenges included:

- Slow searches across large, unstructured datasets
- Difficulty filtering documents by jurisdiction, topic, or case type
- Inconsistent relevance that reduced user trust
- Rising operational costs tied to legacy tools

Deploying Elasticsearch to power fast, precise legal discovery

Legora implemented Elasticsearch to unify structured and unstructured content, enabling hybrid search, filtering, and semantic relevance. The platform delivered faster retrieval, improved accuracy, and scalable performance across an expanding knowledge base.

Achieving faster insights, higher accuracy, and improved user

Legora reduced search latency dramatically and improved relevance across millions of documents. Users now find answers in fewer queries, boosting satisfaction and engagement, and Legora has lowered its operational overhead.

LEGORA

“Elasticsearch gave us the stability, performance, and scalability we required, allowing us to focus on building an innovative product instead of managing search infrastructure.”

Sigge Labor, Co-Founder and CTO, Legora

Empower your organization with lightning-fast, trusted, AI-ready search

Find Elasticsearch in Microsoft Marketplace and accelerate your AI journey.

[Get started here](#)