

	Best for	Integrations	Security/compliance	Pricing	Free trial
Meilisearch	Fast hybrid full-text and vector search; prototyping and production web apps	REST API, 10 language SDKs, OpenAI and Hugging Face embedders	Enterprise: SOC 2 Type II, SSO SAML, RBAC, and audit logs	From \$20/month (cloud); free to self-host under the MIT license	14-day free trial, no credit card required
LangChain	Prototype-to-production RAG pipelines with orchestration and agent flexibility	1,000+ integrations: models, tools, and databases	Enterprise: custom SSO, RBAC	Framework free (MIT); LangSmith Plus from \$39/seat/month	Free Developer plan (one seat, 5k base traces/month)
RAGatouille	Domain-specific retrieval with late interaction scoring; fine-tuning on limited data	LlamaIndex; DSPy	-	Not disclosed	Open source
RAGFlow	Deep document understanding for complex formats: scanned PDFs, slides, tables	OpenAI-compatible API, Elasticsearch, Langfuse, MCP	Admin UI and CLI with user management and system logs; no SOC 2 or HIPAA certification documented as of June 2026	Free to self-host; cloud free tier + paid plans from \$29/month	Free tier available
Haystack	Enterprise RAG pipelines with governance, observability, and managed deployment	REST API deployment, YAML pipelines, Haystack Enterprise Platform cloud deployment	Enterprise Platform: RBAC, SOC 2 Type 2	Framework free; Studio: \$0; Enterprise Platform custom pricing, contact required	Free Studio tier and free Enterprise Platform trial available
Embedchain (Mem0)	Rapid prototyping with minimal setup; early-stage proofs of concept	OpenAI, Anthropic, Hugging Face, local models including Mistral and Llama	Mem0 managed platform: SOC 2 Type I certified, self-attested HIPAA-compliant	Free and open source; Mem0 managed platform from \$19/month (Starter); Enterprise pricing on request	Free open-source self-hosted tier; Mem0 managed free (Hobby) tier available
LlamaIndex	Multi-source RAG with composable indexes, routing logic, and hundreds of data connectors	Data connectors via LlamaHub for a large variety of knowledge and data sources	Enterprise plans include SaaS or hybrid cloud deployment	LlamaParse credit-based from \$1.25 per 1,000 credits	Free plan includes 10,000 credits/month
MongoDB Atlas Vector Search	Teams already on Atlas who want vector search without new infrastructure	Native MongoDB ecosystem; integrates with LlamaIndex and LangChain	SOC 2 Type II, ISO/IEC 27001:2022, PCI DSS, HIPAA-compliant configuration	Free tier (512 MB storage, shared resources, free forever); vector search included in Atlas cluster pricing, pay-as-you-go	Free cluster with 512 MB storage and no time limit
Pinecone	Production-scale similarity search without managing vector infrastructure	REST API, gRPC, LangChain and LlamaIndex; available on AWS	SOC 2 Type II, ISO 27001, GDPR, HIPAA add-on on Standard, included on Enterprise	Starter free; Builder \$20/month flat; Standard \$50/month min; Enterprise \$500/month min	Three-week trial on Standard plan with \$300 in credits
Vespa	Petabyte-scale search and recommendation with custom on-node ML ranking	Custom ranking and on-node ML inference	Vespa Cloud: all communication secured with mTLS; control-plane audit logging; role-based access control	Free to self-host (open source); Vespa Cloud billed per vCPU, memory, and storage	Free self-hosted deployment; Vespa Cloud free trial available, no credit card required