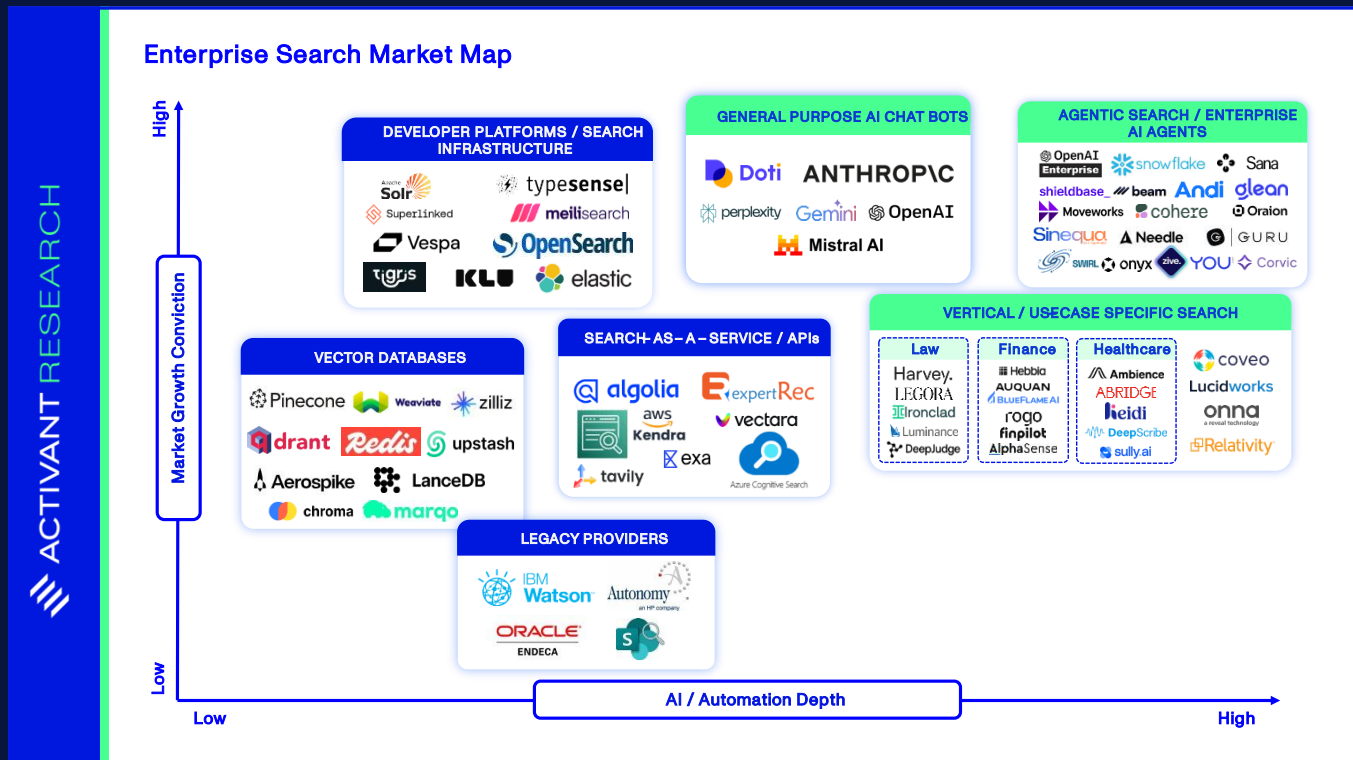




ACTIVANT RESEARCH

Enterprise Search Is Entering a New Era

A shift to systems of intelligence



Rebecca Rodseth, Marc Wu & Teddy Cohen

September 2025

Google gives us an impeccably organized view of the entire internet. Type a query, and the most relevant page appears instantly. Inside a company, however, search often degrades into: “Which folder did I bury that in?” It’s a familiar headache. Now, a new wave of AI-driven search tools is bringing Google-like relevance and speed to our internal knowledge stores.

As AI becomes central to knowledge work, enterprise search is emerging as a critical foundation. Knowledge workers depend on vast unstructured data—across emails, documents, and chats—that traditional search can’t handle. Large language models (LLMs) enable smarter, context-aware discovery, but rely on enterprise search to access and retrieve relevant internal information. Without it, even the best AI operates in a vacuum, disconnected from the knowledge that drives decisions. Enterprise search has expanded beyond internal file systems to include custom applications, infrastructure, and third-party tools. **This new era of information retrieval brings deeper intelligence, reasoning, memory, and agentic research capabilities, enabling smarter, more contextual discovery across a broader data landscape.**

At Activant, we believe that enterprise search could potentially be a game-changer. In line with our work on [systems of intelligence](#), we hold a strong conviction that if all work begins at the enterprise search platform, which centralizes data and systems, it has the potential to become the AI-era “**operating system**” for the modern enterprise.¹

But we’re not there yet. With [79%](#) of employees dissatisfied with internal search functions, the need for enterprise search tools that actually deliver is undeniable.² As companies adopt more SaaS tools (Slack, Notion, Salesforce, Asana, Google Workspace, etc.), information becomes increasingly siloed. As a result, employees spend an average of [1.8 hours](#) a day, or 9.3 hours per week, searching for and gathering information.³ That’s nearly 25% of 40-hour work week lost to inefficiency!

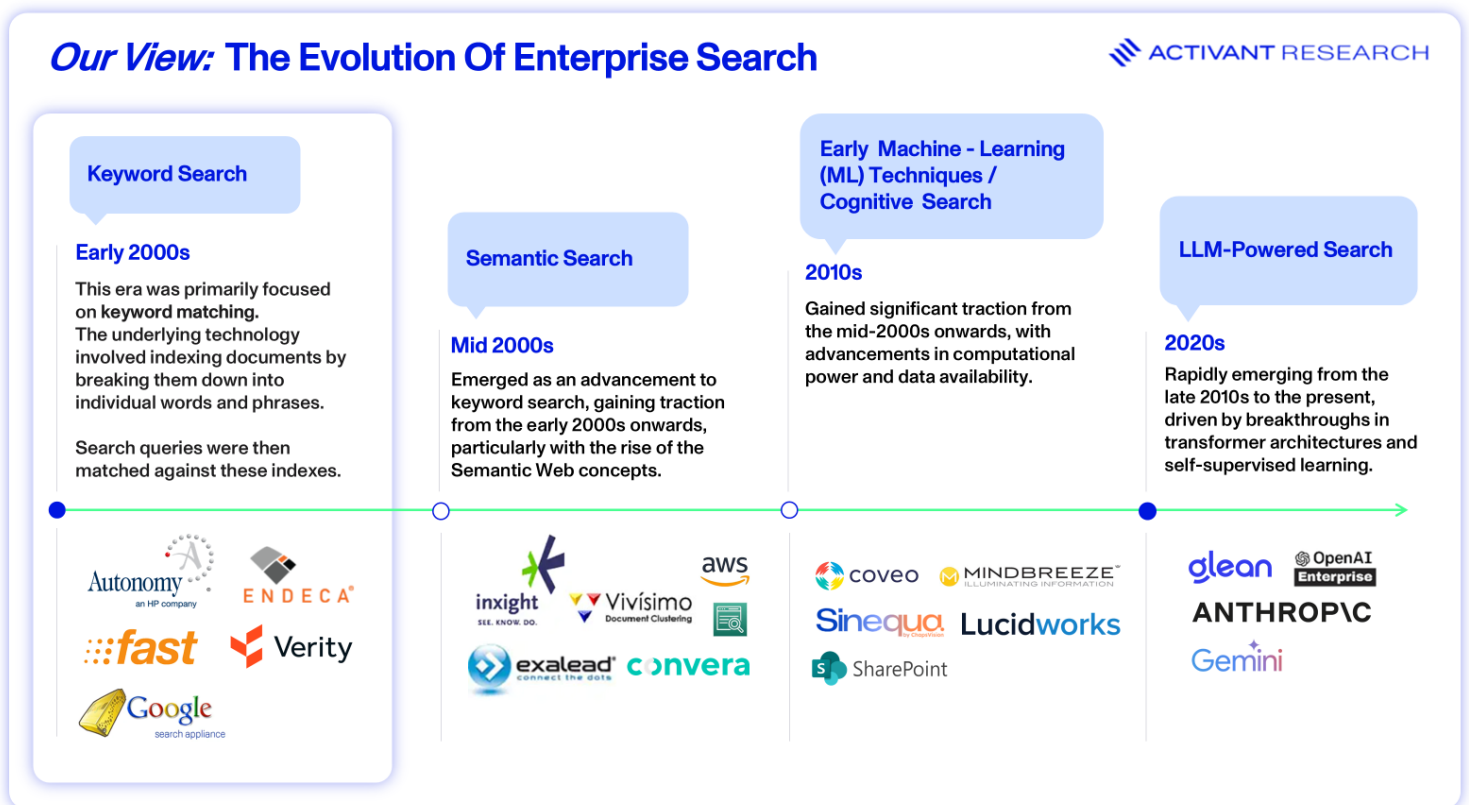
Additionally, horizontal enterprise search solutions have struggled to differentiate meaningfully and often deliver similar generic search experiences across industries and roles. This has caused several early players to rethink their approach. [Genspark](#), for example, publicly scrapped its initial AI search product due to a lack of real traction or unique value.⁴ Similarly, [Glean](#), once positioned as an enterprise search engine, has pivoted toward becoming more of an AI workflow and agent builder, with search as a foundational layer rather than the core product. These shifts reflect a broader realization: **search alone isn’t enough. Successful tools must deliver role-specific utility, context-awareness, and actionability to stand out.**

The new “search war” reveals a deeper paradigm shift: the convergence of external and internal search. Users increasingly expect AI search to not only summarize several sources, but also to intelligently cross-check findings across internal and external sources into decision-ready intelligence. As internal search becomes a standard feature in enterprise platforms and search monetization models evolve, fundamental questions arise about the web’s infrastructure, originally built on clicks, ads, and content distribution. **The agent-native search and agentic web era may be closer than we think.**

Playing Catch Up

You can think of enterprise search as a tool that lets employees find information across all company data sources, like documents, emails, and databases, through one single, unified search interface. Enterprise search isn't new—these tools have been around since the early 2000s. However, their evolution in recent years has been asymptotic, continually approaching but never quite reaching the level of accuracy and usability that's really needed.

The timeline below highlights key technological changes and the relevant impact on the enterprise search industry. It's no secret that enterprise search has failed in the past to deliver against expectations. Early enterprise search solutions were primarily focused on keyword-based searching which often yielded inconsistent results and required manual tuning to improve relevance. However, recent advancements in natural language processing (NLP), machine learning (ML), and artificial intelligence (AI) are improving the capabilities of enterprise search systems. These modern solutions infer user intent and interpret complex queries to deliver accurate, real-time search results. This enhances user experience and drives greater adoption within organizations.



Keyword search laid the groundwork for all subsequent technologies by indexing document content, retrieving results matching exact query terms, and ranking them by keyword frequency or proximity. Simple and inexpensive, it struggled with complex queries or longer phrases.

Semantic search followed and began modelling relationships between concepts using structured vocabularies (ontologies and taxonomies) and knowledge graphs, enabling automatic identification of entities. Although semantic systems improved relevance and data organization, their reliance on manual ontology maintenance limited scalability and rigid predefined concepts restricted their ability to grasp nuanced language or user intent.

Early machine learning techniques automated search by fine-tuning results from user interactions, offering personalized rankings, easier scaling, and improved features like spell-correction and intent detection. However, reliance on handcrafted features, opaque models, costly labelled data, and limited language nuance restricted their effectiveness with complex queries.

The current era focuses on **transformer-based LLMs** trained on vast text data, that group semantically similar content and generate fluent, context-aware outputs. Enterprises increasingly pair tool-using LLMs with multi-modal databases for accurate, specialized results but challenges like compute costs, latency and data security persist.

Up to [80%](#) of enterprise data is unstructured. Modern enterprises are drowning in the vast amounts of data that they are generating and storing.⁵ This includes emails, pdfs, chats and internal wikis. The recent addition of AI to enterprise search tools has created capabilities like semantic search and various retrieval strategies. This allows users to get direct answers, not just documents, grounded in real-time organizational knowledge. In short, as data sprawl collides with LLM-driven user expectations, enterprise search is evolving from a “nice-to-have” lookup tool into the critical intelligence layer that lets companies capitalize on the information they already own. **With the increased ability to deal with complexity, LLMs mean that search shifts from just finding the data to starting to perform the actual work for you.**



While investment in data infrastructure is at an all-time high, enterprises face a persistent talent shortfall. This gap in data expertise delays critical AI and search deployments, leaving vast stores of unstructured information underutilized.”

Farshid Sabet

Co-Founder & CEO, Corvic AI

Generative AI Is Addressing the Challenges of the Past

AI is the major change that has unlocked the next level of search. As AI becomes more integrated within enterprise search tools, many of the traditional challenges are being addressed:

- **Data silo integration:** Enterprises store information in countless systems. As a result, [87%](#) of organizations struggle with disconnected data sources, leading to inefficiencies in operations and decision-making.⁶ In the past, a major challenge has been connecting to and indexing all of these disparate sources so that search truly becomes universal. Modern retrieval frameworks now embed LLMs that auto-detect file types, infer schemas and normalize metadata as content is ingested.
- **Security, privacy and compliance:** Indexing or embedding sensitive documents may violate General Data Protection Regulation (GDPR) and California Consumer Privacy Act (CCPA) or leak contract terms if access controls aren't honored end-to-end. The use of LLMs has enabled enterprise search players to implement permissions when collecting from source applications. Encryption, row-level permission sync and audit trails are now table stakes but they are expensive to retrofit. Another option is private/on-prem LLM deployments that keep embeddings and inference fully behind the firewall for organizations that can't send data to public clouds.
- **Scalability and performance:** AI has made enterprise search more scalable and performant by replacing rigid keyword-based systems with models that understand meaning across large, messy datasets. It reduces manual setup, compresses search space semantically, and intelligently routes queries—delivering faster, more accurate results even as data and user demand scale.

As data volume explodes, scaling a search solution can be challenging both technically and financially. Data growth means continually expanding indexes, which impacts storage costs and query performance. Specialized vector databases are able to deliver millisecond queries at billion-vector scale, with storage-optimized tiers that cut costs by a factor of [5-7x](#).⁷ Given that [90%](#) of users expect responses to show in under 2 seconds, failing to meet this modern benchmark hinders the adoption of enterprise search tools.⁸

- **Hallucinations and answer fidelity:** Early LLMs functioned primarily as sophisticated text predictors, often struggling with factual reliability due to limited reasoning capabilities. Initial AI search products like [Perplexity](#) addressed this by manually orchestrating query decomposition and keyword-based API searches, then using LLMs to compile the results.

The landscape has since shifted with the emergence of advanced reasoning models that can break down complex queries and autonomously select the right tools or data sources. This evolution has significantly improved factual accuracy and reduced the reliance on rigid retrieval frameworks. Platforms like [ChatGPT](#) now dynamically integrate real-time search and data-grounding mechanisms, rather than relying solely

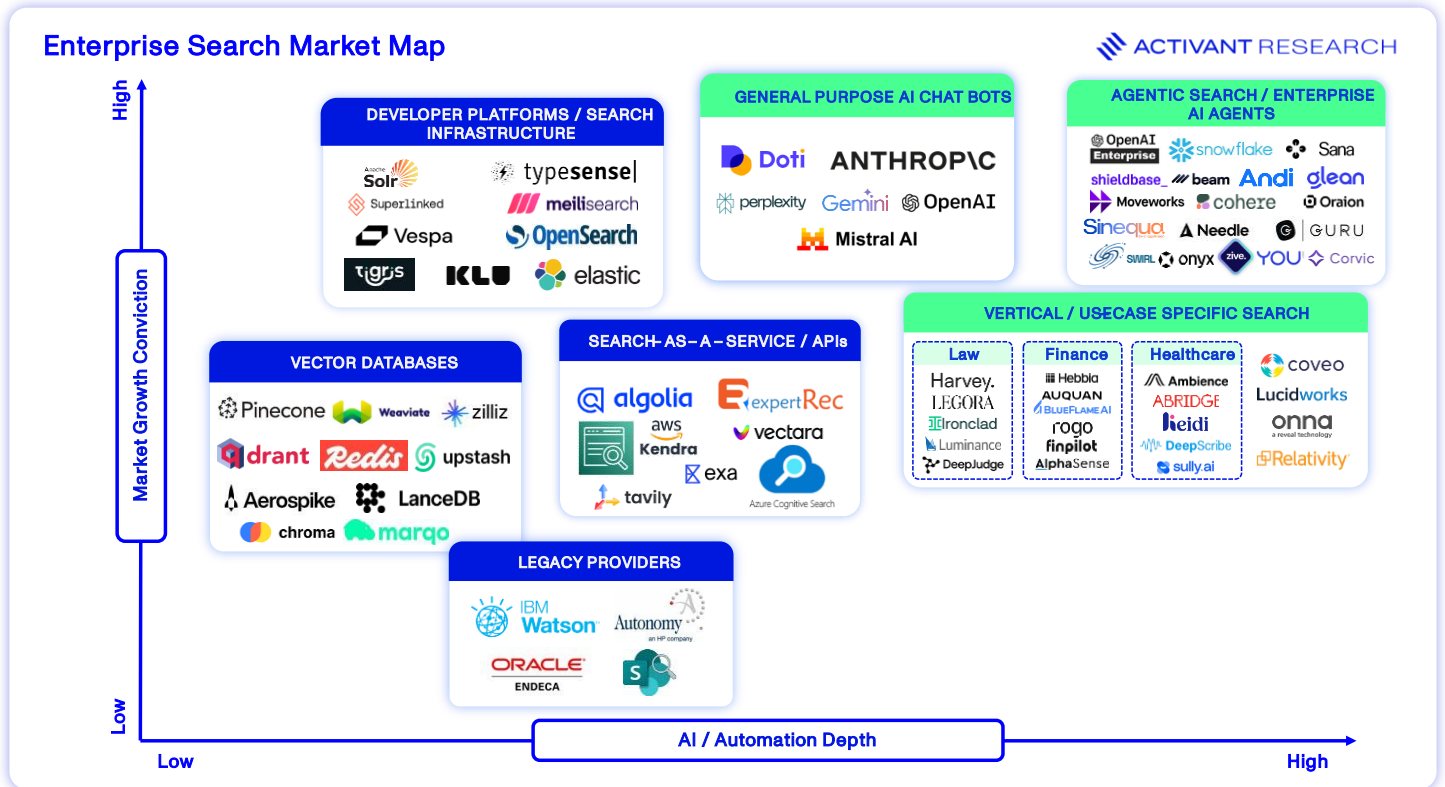
on pre-encoded knowledge or static retrieval. This approach delivers high-precision answers with significantly fewer reliability concerns.

- **Relevance tuning:** Implementing search isn't just a one-time technical task; tuning relevance to satisfy users is an iterative and ongoing challenge. Relevance tuning involves making sure that search results are the most relevant ones and also the most recent—users want the latest information available. Out-of-the-box relevance may be inadequate because every organization has a unique mix of content and jargon.

The improved capabilities of search tools have attracted new entrants to the market, as well as upgraded the offerings of existing legacy players.

How We See the Market

Our Activant market map assesses companies based on our market growth conviction and perception of the level of AI / Automation Depth. The rating of AI / Automation depth attempts to capture how much cognitive work a search product assumes, ranging from basic keyword matching at the bottom to fully autonomous, LLM-driven agents at the top. We are most excited about the opportunities within **agentic search** and **vertical/ use-case specific search**. There are naturally opportunities across the market, and one cannot ignore the fact that **general purpose AI chat bot** providers such as Google, OpenAI and Anthropic are well-positioned to own enterprise search if they choose to invest appropriately.



Agentic Search/Enterprise AI Agents leverage advanced natural language processing (NLP), reasoning capabilities, and autonomous decision-making to act as intelligent agents. Rather than simply retrieving links, they interpret queries contextually, synthesize information, and deliver concise, tailored responses. Think of them as digital assistants with the ability to reason, adapt, and even anticipate user needs, far beyond the static retrieval of Google’s keyword search model.

The rise of agentic search turns enterprise search from a passive lookup tool into the cognitive motor of autonomous workflows. Vendors that can feed, steer, and govern agents, and enterprises that modernize their search stack accordingly, will capture outsized productivity and insight gains over the next five years.

We are excited about [Sana](#), a no-code cloud platform that allows organizations to deploy AI agents grounded in their internal knowledge and tools. [Onyx](#) is an open-source platform that turns a company’s scattered docs, apps and messages into a single AI-powered “coworker.” It syncs permissions across 40+ connectors such as Google Drive, Slack, Confluence and Salesforce. [Oraion](#) is innovating at the intersection of enterprise intelligence, automation, and agentic AI, a step beyond traditional analytics and conversational bots. It positions itself as a “single source of truth,” intelligently connecting and interpreting data from over 300 sources, via a natural-language interface.



Evolving from enterprise search into an AI Operating System is not only about technology. It requires rethinking workflows around AI-augmented processes, ensuring quality data with robust governance, while also upskilling the workforce to collaborate effectively with intelligent systems.”

Diego Rojas

Founder & CEO, Shieldbase AI

Vertical/use-case specific search includes vendors that pre-train models on a profession’s data and workflows, so relevance, compliance rules, and UX arrive out-of-the-box. The tight domain focus commands premium pricing and usually beats horizontal engines on recall and speed for that niche.

We’re bullish on this category and our previous work on [vertical software](#) has laid a strong foundation for where we see the market heading.⁹ Vertical software vendors have started baking enterprise-search directly into the workflows they already own. This is a strong value proposition. These platforms operate directly within the daily workflows of analysts, lawyers, or product managers, inheriting rich context such as metadata, permissions, and task states that horizontal engines lack. This integration surfaces insights precisely at decision points, eliminating the need for users to switch portals.

Looking at a player like [AlphaSense](#) shows how this plays out in financial research. Its platform already streams broker reports, earnings calls, and proprietary notes into analysts’ dashboards. Users stay in AlphaSense, click once, and see cited passages ready for modelling, without context-switching to a generic search box. [Hebbia’s](#) core product, Matrix, lets knowledge-heavy teams ask natural-language questions across files. The system returns spreadsheet-like answers with linked source citations. It can also extend into multi-step AI agents for tasks such as due-diligence data extraction, memo drafting and valuation modelling. [DeepJudge](#) is an AI platform tailored for legal professionals. At its core lies DeepJudge Knowledge Search, an enterprise-grade, AI-powered engine that indexes information across document management system (DMS) platforms, SharePoint, emails, intranets, and more. It retains all access controls, ensuring lawyers only see what they’re authorized to view.

These examples underline a broader shift: **vertical software players aren’t just integrating search for convenience—they’re turning it into a domain-aware co-pilot** that accelerates core tasks, giving them both stickier products and a new, AI-driven upsell path.

General-purpose AI chat bots: These are essentially foundation-model providers that also offer enterprise search, not pure-play search startups. So, while specialist firms like Perplexity were “search-first” from day one, the big-model vendors like OpenAI, Google, and Anthropic have crossed the same threshold by packaging vector retrieval, permission handling and answer generation around their LLMs. **Horizontal search can index everything, but depth beats**

breadth when the business value hinges on finding exactly the right answer. The fact remains that without enterprise search capabilities, LLMs will never be reliable enough for large scale automation of knowledge work.

ChatGPT Connectors is already live in beta for business users. Originally announced in March 2025, the feature allows ChatGPT Team subscribers to securely connect platforms like Google Drive and Slack, enabling the AI to search internal documents and messages. Since then, the functionality has expanded to include integrations with Dropbox, SharePoint and OneDrive, and is available to Team, Enterprise, and Education-tier users.

In December 2024, Google unveiled Google Agentspace. This layers Gemini reasoning over a unified knowledge graph so employees can query and act on data across Workspace, Microsoft 365, Salesforce, ServiceNow and other systems from a single prompt. In addition, Google's Vertex AI Search gives developers Google-quality retrieval and retrieval-augmented generation (RAG) with ready-made connectors to Salesforce, Jira, Confluence and more.

Anthropic moved into enterprise search with Claude Enterprise in September 2024. Claude can query private code and documents with enterprise-grade controls. In April 2025, it added an agentic research mode and a Google Workspace connector that enables Claude to perform permission-aware searches across Gmail, Calendar, and Drive. These results are blended with live web retrieval to return cited answers, directly rivalling Copilot-style workplace search assistants. In July 2025, Anthropic also launched **Claude for Financial Services**, a purpose-built version of its Claude 4 AI models tailored specifically for the financial sector. It integrates seamlessly with platforms like Snowflake, Databricks, PitchBook, Morningstar, S&P Global, FactSet and more, offering a unified interface that connects internal and external data via secure MCP connectors.

Search-as-a-Service/APIs are ready-to-use, cloud-based tools that handle all backend operations. You simply provide your data and receive a reliable, scalable search or RAG API in return—billed by usage, with the vendor managing uptime, security, and system tuning.

[Vectara](#) offers an API-first, “RAG-as-a-Service” platform. Developers stream their data in, and it handles embeddings, retrieval, grounded answer generation, and now autonomous “guardian” agents to detect and even *correct* hallucinations. We are also interested in [Exa](#), which bills itself as “the search engine for AIs.” With AI companies now its primary customers, use cases for Exa’s search engine span everything from an AI chatbot looking up info on the internet while answering customers’ questions, to companies looking to curate training data.

Developer platforms and search-infrastructure projects provide search engine software to run yourself, either in your own cloud environment or on your company’s servers. You’re in charge of setting it up, managing it, and making any changes. They’re a good fit for teams that need full control over how search works or where data is stored. But they also require your team to handle all the maintenance, updates, and troubleshooting.

We’re excited about [Meilisearch](#), which has developed a lightning-fast, open-source search engine written in Rust and offered under the MIT license. The project has become a developer

favorite, passing 52,000 GitHub stars in 2025. [Klu](#) is a generative AI platform that simplifies the process of designing, deploying, and optimizing AI applications.

Vector databases have previously been the backbone of enterprise-grade search and generative-AI projects. This was because of their ability to store data as high-dimensional embeddings, allowing systems to retrieve information by meaning rather than by exact keywords. When a language model needs to ground its answers in company knowledge (the RAG, pattern), it first looks up the most relevant vectors, speed and recall that dedicated clouds deliver at millisecond latency even at billions of records.

However, it's important to flag that users find that standalone vector databases are not enough. Instead, customers prefer multi-modal databases. Organizations face scalability and cost challenges, and newer technologies—such as hybrid search, graph databases, and advanced indexing—offer compelling alternatives. As a result, enterprises now prefer a more diverse set of tools to meet evolving AI demands.

We like [Marqo](#), a vector search platform that integrates applications, websites and workflows. Managing the entire process from vector creation to storage and retrieval, Marqo facilitates the effortless integration of multimodal, multilingual search via a unified API. [Qdrant](#) builds an open-source vector database and similarity-search engine written in Rust. Designed for RAG and other AI workloads, it supports hybrid filtering, any text or multimodal embeddings, and scales to billions of vectors via gRPC/REST APIs. [Zilliz](#) builds the open-source Milvus vector-database engine and sells Zilliz Cloud, a fully managed, serverless version that runs on AWS, Azure and GCP.

Legacy providers are long-standing suites that sprung from keyword-based insight engines but now bolt on vector search and GenAI plugins. They still win with huge connector catalogs, mature governance features, and on-prem options, making them the default for compliance-heavy or already-installed enterprises.

It's evident that organizations have a plethora of enterprise search providers to choose from. This is a competitive space. Only those that are on top of their game and are customer-centric will be able to gain real market share.

Key User Needs

Employees often run into issues when trying to use internal search tools. Users want instant, trustworthy answers from the tools they already use, whether that means a chat prompt that cites sources or a results page that filters by project and freshness. As a result, we are seeing a huge demand for both deep research and agentic search. Players should incorporate these key user needs:

- **Concurrent internal and external search** where systems intelligently alternate between internal and external searches. Solutions addressing this emerging need leverage advanced “deep research” systems enabled by improved reasoning models, multi-

agent architectures, and communication protocols like Model-Context Protocol (MCP) to perform up to 100 times more iterations for greater precision, albeit at higher cost. There is also a shift toward broader “dataset” searches where multiple agents simultaneously explore internal and external data in parallel. This evolution underscores a crucial transition toward asynchronous AI experiences, allowing agents to manage complex, multi-step workflows independently, delivering curated outcomes without interrupting user interactions.

- **A single search interface** that offers unified data integration, eliminating silos and saving users from hunting in multiple places. Today, [81%](#) of organizations report that data silos are hindering their digital transformation efforts.¹⁰ Emerging standards like MCP are helping to address this challenge by enabling AI systems to interface with disparate tools and data sources through a unified, extensible protocol. Rather than relying on brittle, hardcoded integrations, MCP facilitates dynamic tool use and retrieval across organizational systems. This makes unified search not only possible, but efficient, secure, and adaptable at scale.
- **Semantic relevance and natural language understanding** to interpret conversational queries, synonyms and domain-specific jargon rather than relying on exact keyword syntax. This allows users to find answers based on what they mean, not just what they type. It's important for enterprise systems to extract information not just from unstructured data, but also from structured sources using methods like text-to-SQL. This kind of intelligent auto-routing across both data types enables more comprehensive and accurate information retrieval.
- **Intuitive, customizable UX** design to increase adoption and user satisfaction. Effective interfaces support features like faceted search, filters, and sorting options to help users quickly narrow down results. Faceting allows users to sort results based on specified fields. Users typically adapt to the designs of tools that they encounter first. As a result, many enterprise search players have mirrored ChatGPT's UI/UX design. But with the rise of agent capabilities, we're now seeing a shift towards making AI assistants the primary interface.
- **Robust security and compliance** to address privacy concerns. Cisco notes that [27%](#) of organizations have banned GenAI tools over data security risks.¹¹ Regulatory differences across regions heavily influence market dynamics, often determining vendor dominance more than technological capability alone. Stringent regulations like GDPR favor vendors with robust privacy and data residency capabilities in Europe, while U.S.-based vendors benefit from less centralized regulatory environments in North America. In China, unique data localization and censorship laws restrict market entry, favoring domestic players like Baidu. Consequently, vendor success frequently

depends on alignment with regional legal frameworks and cultural expectations rather than purely technological strengths.

- **Deployment flexibility.** Organizations require flexible enterprise search solutions tailored to their infrastructure needs, with options for cloud, on-premises, or hybrid deployments. Cloud solutions offer ease of maintenance and automatic updates, while on-premises setups provide superior control over data security. Increasingly, hybrid models allow organizations to balance these advantages. Although traditionally challenging, on-premises deployments are becoming more accessible with custom offerings from providers like [Glean](#).
- **Continuous learning and analytics** that help organizations understand how users interact with their search functionality. These insights drive improvements in search relevance and user experience. Enterprise search platforms collect and analyze data about search patterns, popular queries and failed searches. Modern analytics tools go beyond basic metrics, offering AI-driven insights that identify trends and patterns in search behavior. This information helps organizations optimize their content strategy and improve search effectiveness over time.
- **Context and Memory** that delivers results tailored to a user's role, past behavior and current project context (e.g., team, department or language) so every search feels like a personal assistant rather than a generic engine. In April 2025, OpenAI expanded memory capacity by ~25% for paid tiers, enhancing ChatGPT's capacity to retain and recall across sessions.¹² Context engineering has taken off as a crucial skill in building reliable, agent-based AI systems. LangChain recently described this as "the art and science of filling the context window with just the right information for the next step," a shift driven by the rise of multi-step agents that rely on dynamic context management, not just static prompts.¹³

It's all good and well winning the customer. But questions around the monetization of these tools remain.

Monetizing Enterprise Search

As enterprise search platforms evolve from passive query tools into dynamic AI agents, their monetization models are rapidly shifting to reflect the explosion in underlying compute demands. Historically, vendors charged a flat per-user fee which aligned with predictable, human-initiated search behavior. But the rise of agent-based search, where AI agents run multiple complex searches at once across large sets of information, has significantly increased backend workload and token consumption.

As a result, many enterprise search companies are moving toward hybrid or fully usage-based pricing, often centered around tokens or query volume thresholds. This model better aligns cost with usage intensity: light users remain low-cost, while high-throughput agent interactions incur overage pricing. Like API pricing in LLMs or observability tools, this shift captures the value of real-time inference. It also creates a usage-based incentive, encouraging efficient orchestration and pushing buyers to closely monitor search agent sprawl. Going forward, we may also see vendors introduce tiered pricing tied to latency, model size, or vertical-specific connectors, further segmenting value in an increasingly intelligent enterprise search stack.

As a result of the shift from static query tools to dynamic, agent-driven systems, the *“economics underpinning enterprise search platforms is undergoing a fundamental transformation.”* As enterprise search evolves into dynamic, LLM-based agent systems, query economics has shifted dramatically. Traditional web searches like Google’s cost about \$0.0106 per query in operational overhead and yield roughly \$0.0161 revenue per search.¹⁴ By contrast, GPT-5 models charge around \$0.000125 per 1,000 input tokens and \$0.001 per 1,000 output tokens, meaning a typical 1,000-token query-and-response costs approximately \$0.001125, or 11¢.¹⁵ **This disparity—11¢ per agent query versus ~1¢ for basic search—underscores why flat per-seat pricing no longer aligns with usage.**

The recent breakthroughs in deep research and parallel asynchronous search may further exacerbate the problem. Each deep research query performs 10–100x more tool calls and associated token usage, which translates to an equivalent price increase. If we assume that every agent in a multi-agent framework performs such lengthy steps, the ensemble approach may then hike the price 5-10x, resulting in ~1,000x price increase on inference cost alone. Fortunately, LLM inference price is also dropping quickly. [Based on Epoch's data insights](#), GenAI inference cost decreases range from 9 -900x per year across tasks and benchmarks, partially offsetting the new compute intensity.¹⁶ Monetization challenges on the revenue side, however, still exist.

An Exciting Future Awaits

The way we access knowledge at work is undergoing a transformation. While Google offers a seamless experience navigating the vastness of the internet, internal company search remains fragmented and frustrating. AI is changing that. The rise of intelligent search tools brings newfound precision and speed to enterprise knowledge, promising to eliminate the daily grind of information hunting.

At Activant, we see enterprise search as a potential cornerstone of the modern enterprise, an AI-first interface that unifies systems, orchestrates workflows, and becomes the default entry point to all work. But there’s a gap between that vision and today’s reality. With most employees unhappy with internal search and valuable time lost daily in information silos, the urgency is clear.

Generic, one-size-fits-all tools have failed to meet the mark. The next wave must offer tailored insights, adapt to different roles, and deliver tangible value beyond static queries. As startups pivot from pure search toward integrated AI platforms, it's clear that the future lies in action-oriented intelligence, not just information retrieval. Ultimately, we're seeing a shift, where **search is no longer just about finding, but about understanding, connecting, and deciding**. The boundary between internal and external search is fading, and with it, the foundations of the web itself may evolve. In this emerging era, enterprise search won't just support work, it may very well define it.

We would like to hear from you if our work resonates or if you have a different perspective. If you're building in this space, we'd love to connect!

-
- ¹ Activant Research, [Systems of Intelligence](#), 2025
 - ² Zive, [Enterprise Search](#), 2024
 - ³ IBM, [IBM Support](#), 2025
 - ⁴ MainFunc, [Why I killed Our AI Search Product With 5 Million Users](#), 2025
 - ⁵ Bizdata, [Structured vs Unstructured Data: Comprehensive Guide](#), 2025
 - ⁶ Bizdata, [What are Data Silos? Problems and Solutions Guide](#), 2025
 - ⁷ Databricks, [Announcing Storage – Optimized Endpoints For Vector Search](#), 2025
 - ⁸ Guru Software, [The Complete Guide To Response Time Testing And Optimization](#), 2024
 - ⁹ Activant Research, [Vertical Software Is Having A Moment](#), 2025
 - ¹⁰ Salesforce, [85% of IT Leaders See AI Boosting Productivity, but Data Integration and Overwhelmed Teams Hinder Success](#), 2024
 - ¹¹ Cisco, [More than 1 in 4 Organizations Banned Use of GenAI Over Privacy and Data Security Risks](#), 2024
 - ¹² OpenTools, [OpenAI Expands ChatGPT Memory by 25% for Premium Users](#), 2025
 - ¹³ LangChain, [Context Engineering](#), 2025
 - ¹⁴ Semi analysis, [The Inference Cost of Search Disruption – Large Language Model Cost Analysis](#), 2023
 - ¹⁵ Partenit, [The Cost Equation: How Ontology Queries Slash Your GPT Token Budget By 90%](#), 2024
 - ¹⁶ Epoch AI, [LLM inference prices have fallen rapidly but unequally across tasks](#), 2025

Disclaimer: The information contained herein is provided for informational purposes only and should not be construed as investment advice. The opinions, views, forecasts, performance, estimates, etc. expressed herein are subject to change without notice. Certain statements contained herein reflect the subjective views and opinions of Activant. Past performance is not indicative of future results. No representation is made that any investment will or is likely to achieve its objectives. All investments involve risk and may result in loss. This newsletter does not constitute an offer to sell or a solicitation of an offer to buy any security. Activant does not provide tax or legal advice and you are encouraged to seek the advice of a tax or legal professional regarding your individual circumstances.

This content may not under any circumstances be relied upon when making a decision to invest in any fund or investment, including those managed by Activant. Certain information contained in here has been obtained from third-party sources, including from portfolio companies of funds managed by Activant. While taken from sources believed to be reliable, Activant has not independently verified such information and makes no representations about the current or enduring accuracy of the information or its appropriateness for a given situation.

Activant does not solicit or make its services available to the public. The content provided herein may include information regarding past and/or present portfolio companies or investments managed by Activant, its affiliates and/or personnel. References to specific companies are for illustrative purposes only and do not necessarily reflect Activant investments. It should not be assumed that investments made in the future will have similar characteristics. Please see “full list of investments” at <https://activantcapital.com/companies/> for a full list of investments. Any portfolio companies discussed herein should not be assumed to have been profitable. Certain information herein constitutes “forward-looking statements.” All forward-looking statements represent only the intent and belief of Activant as of the date such statements were made. None of Activant or any of its affiliates (i) assumes any responsibility for the accuracy and completeness of any forward-looking statements or (ii) undertakes any obligation to disseminate any updates or revisions to any forward-looking statement contained herein to reflect any change in their expectation with regard thereto or any change in events, conditions or circumstances on which any such statement is based. Due to various risks and uncertainties, actual events or results may differ materially from those reflected or contemplated in such forward-looking statements.