


Enterprise GenAI: To Build or Buy a GenAI LLM?



AlphaSense

What's Inside

1 INTRODUCTION

3 PART 1

**The Costs of Inefficient
Knowledge Sharing**

5 PART 2

The Potential of GenAI

6 PART 3

The Risks of GenAI

8 PART 4

To Buy or to Build?

10 PART 5


**End-To-End Solutions
Extend Beyond LLM or
Vector Database Decisions**

12 PART 6

**AlphaSense is Your
AI-Powered Insight Hub**

16 PART 7

**Benefits of AI Generally for
Internal Content**



Introduction

Today, the information challenges organizations face, including leveraging their own existing research and insights, exceeds overload. Nearly every organization has previously undertaken the implementation of a Knowledge Management System (KMS), only to be met with outcomes that fell short of expectations.

This is a reality that many organizations continue to face: their data is a mess and even when they know where to look within a KMS, it takes 20 clicks in folders to access it, and then relying on CTRL+F to search a tidbit of information within a document. Content finds itself in various locations, whether on a local drive, within a team folder, or within a team-centric note-taking application, excluding broader organizational access.

This roadblock can escalate into a formidable complexity, particularly when managing separate investment teams, conglomerates with diverse business lines, or regional teams. It also leads to teams working with a lot of mismanaged internal and external information that wastes money and time.

Furthermore, invaluable intellectual property inherited from previous investments, projects, deals or acquisitions often remains ensnared within legacy systems, separate and detached from the central knowledge repository.

These challenges lead to a reliance on multiple different storage platforms and missed information. The result of these costly mistakes? Laborious data gathering that is too slow for the pace of the business, endless duplication of efforts across different teams, and diminished operational efficiency on research. Moreover, it leads to misinformed decision-making with significant downstream impacts to the organization.

In terms of downstream impacts, we're talking about wasted money, billions of dollars annually on sourcing research content, and staffing costs related to building internal research. So what becomes the solution to navigating the fury of external content alongside the endless reservoirs of internal insights that often proves to be even more valuable?

The Costs of Inefficient Knowledge Sharing

The challenges of sharing internal information carry substantial repercussions when we consider the tangible financial impact stemming from the cumulative effects of wasted time and suboptimal decision-making.

According to assessments by [Panopto and YOU.gov](#), major U.S. corporations suffer annual losses exceeding \$40 million as a result of everyday operational inefficiencies directly linked to inadequate knowledge sharing.

Lost Productivity

\$40mm+

Average productivity loss for large U.S. companies

Missed Opportunities

5.3 hours

Wasted each week per employee waiting for information

Knowledge Loss

200+ hours

Time inefficiency for average new hires

Sources: Panopto, HR Daily Advisor

[Gartner](#), likewise, posits that suboptimal operational decision-making inflicts a cost equivalent to 3% or more of a company's annual profits. Based on the data furnished by the St. Louis Federal Reserve regarding corporate profits in 2022, this translates to a staggering loss exceeding \$350 billion within the United States alone during the preceding year.

Additionally, Panopto's research also reveals a concerning trend—[new hires typically spend over 200 hours on inefficient activities](#) as they familiarize themselves with their new roles, largely due to shortcomings in knowledge transfer and the accessibility of vital information.

By translating the challenges to the bottom-line impacts, you can understand why companies are looking for solutions to help them drive faster, more efficient knowledge sharing at their organizations. And the seemingly obvious solution to these challenges is better technology that helps surface relevant information quickly.

Cue genAI.



The Potential of GenAI

From automating data analysis and forecasting to generating personalized investment recommendations, detecting patterns and building out prediction models, the adoption of AI tools is not just a trend, but a strategic move that can drive innovation, operational efficiency, and success

Some market intelligence platforms leveraging genAI unravel the intricacies of unstructured data, to promptly extract valuable insights from the ever-expanding lake of institutional intellectual property. The goal is to rejuvenate existing content repositories, effectively breathing new life into documents, and extracting fresh value to enrich contemporary decision-making processes.

Today, it seems professionals amongst industries are facing a similar need: to render internal research as discoverable and manageable as external content that many vendors provide the marketplace. So how can internal research achieve this element of being “clean, accessible” data on an enterprise scale?

The answer lies in harnessing the power of artificial intelligence and machine learning. By comprehensively understanding not only the content of each document but also the business context of each paragraph, each sentence, each word used; this technology makes sense of data in ways that professionals never thought possible.

In our [2023 State of Gen AI & Market Intelligence report](#)—which surveyed 500-plus professionals across various roles—a vast majority (over 80%) of respondents plan to leverage genAI tools in their research this coming year.

The Risks of GenAI

While many are rushing to tout genAI as the solution to workplace obstacles and streamlining process, some are questioning its trustworthiness, data security, and ease of implementation. Organizations are trying to understand what is the investment that will be required and if the risks are worth the effort.

When it comes to expectations for genAI versus the reality of this technology, there's a major discrepancy between the two in the eyes of professionals. While some have lofty visions of building a genAI that can easily be converted into code and algorithms—and with this mindset, rush to market with constructing a tool, this leads to a few potential problems.

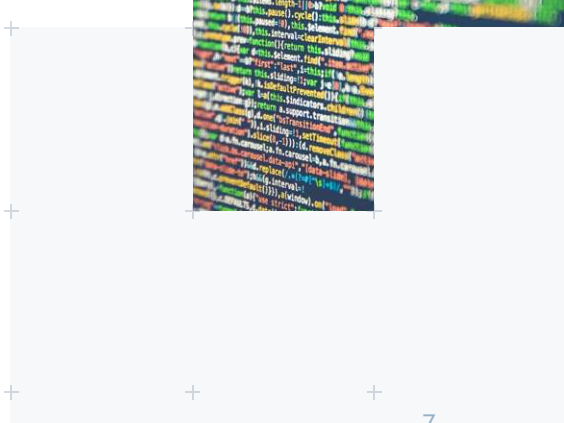
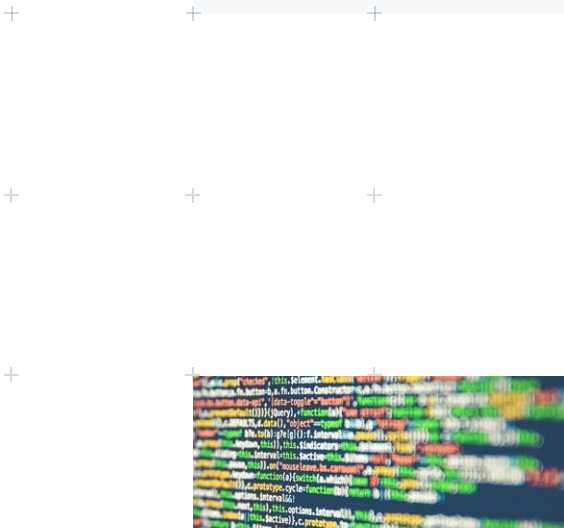
First, a genAI tool and its algorithms could be trained on misleading or inaccurate information that is based on the entire internet or non-authoritative sources. This leads to it generating outputs that are difficult to trust without knowing the origin of information, making it challenging to assess the bias of the original source or citation. Ultimately, these non-vetted or untrustworthy tools have a higher likelihood or potential to generate fake facts (i.e hallucinate) while sounding realistic or trustworthy.

In terms of accuracy, many popular genAI models are built for general use and have been trained on data that ranges broadly in quality and relevancy. By relying on unverified content that may propagate biases and models that cannot discern what is misinformation or disinformation, these general models are problematic when applied to business purposes. The generated responses they produce are often lower-quality, or worse yet, completely false but very convincing (otherwise known as “hallucination”).

That's why it's critical to ensure generated responses come with transparent sourcing for validation purposes. But there are also risks around data privacy, as some models may or may not use client data as materials to train. At AlphaSense, we ensure that our proprietary AI algorithms and genAI tools never leverage proprietary client information to train our large language model (LLM).

In the realm of SaaS solutions, security is a paramount consideration. That's why many are contemplating the adoption of closed systems, and focusing on where their data and insights remain within the confines of their organization.

In light of this, organizations, spanning financial services, corporations, and consulting firms, find themselves in a critical deliberation phase: **whether to procure an all-encompassing, off-the-shelf solution or embark on the journey of crafting a bespoke solution from the ground up.** And in other cases, some believe this dilemma is much simpler by attempting to buy an OpenAI license and a Vector Database—leading them to believe they are buying, but in fact, need to build. This strategic dilemma has become a focal point for CTOs, CIOs and IT teams, engaging in robust discussions to determine the best path forward.



To Buy or to Build?

Not only is it one of the most important technology and business decisions to be made in the next 12 months, but it has the potential to also be one of the most expensive. Building LLMs takes significant time, resources, and capital in order to effectively acquire, clean, and curate large data sets. Just for model training, you need computational resources (i.e., hardware, software, cloud services), ramping personnel, maintenance and updates, the list goes on and on.

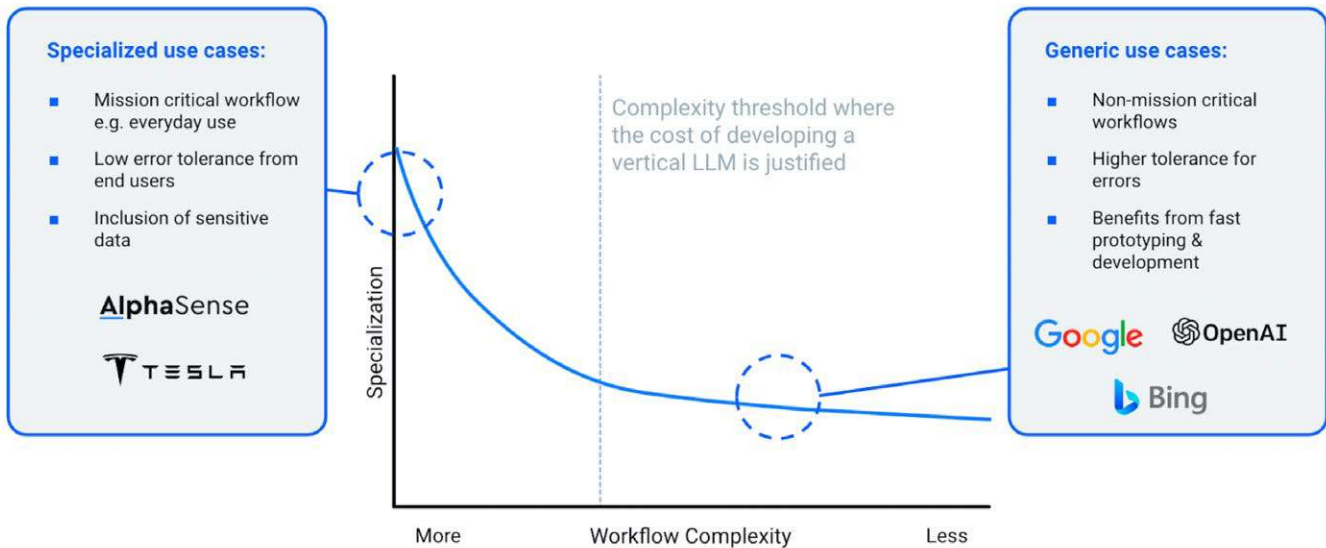
The result of asking these hard questions leaves companies in a tough spot: they realize the potential positive impacts genAI can have on their business but implementing it is challenging. The classic “buy vs build” debate rages on within firms of all shapes and sizes.

Exploring external solutions leads to questions about data security and data control, flexibility and customization, reliability and scalability, implementation times, and implementation costs. So it becomes crucial to understand why and how you’ll be using an LLM.

Determine the Use Case to Drive Gen AI Model Selection

The most important question when it comes to deciding whether to build or buy an LLM is what are you looking to achieve with the solution? Your use case should drive what type of model and overall product your company settles on.

If you intend to use genAI for lower complexity work, like brainstorming positioning statements, email copy generation, or other non-mission critical work, a general-purpose model may be a great fit. However, in the case of strategic work with high complexity, these models are less likely to perform to your advantage.

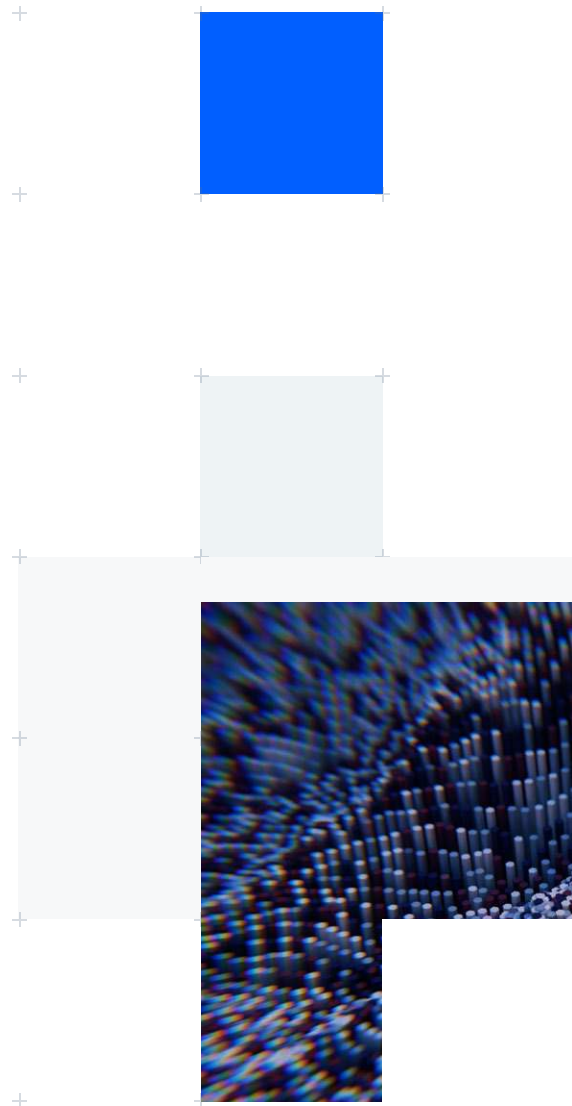


But as the requests you have for genAI become more complex, so does the need for a higher level of accuracy, and the error tolerance decreases. Especially in the case where there is a higher security threshold for source content, more specialized models may better fit a client's requirements.

For example, when considering potential acquisitions, you need to perform comprehensive due diligence that encompasses an extensive overview of the investment landscape, target companies, and competitors. The same is required of any major investment or expansion strategy—the latest, most thorough information required to make a sound, smart business decision.

The greatest impact a large language model can deliver is where it has domain expertise, thus delivering higher precision and recall, improved information retrieval, less bias, computational efficiencies, and targeted knowledge extraction—all leading to an improved user experience. This is especially true for those whose application is in the field of information and knowledge management.

An evaluation of model characteristics could include mapping your use case requirements against factors such as a model's factual consistency, coherence or fluency, granularity of detail, or redundancy in output responses.



End-To-End Solutions Extend Beyond LLM or Vector Database Decisions

After you've evaluated your use case fit, you may want to employ a more specialized base LLM model to meet your organizational needs. In the event you choose to build, you'll need a foundational LLM with a vector database to improve the relevancy of responses, as well as the speed and efficiency of your model. However, that's not the full extent of what you would need to 'build.'

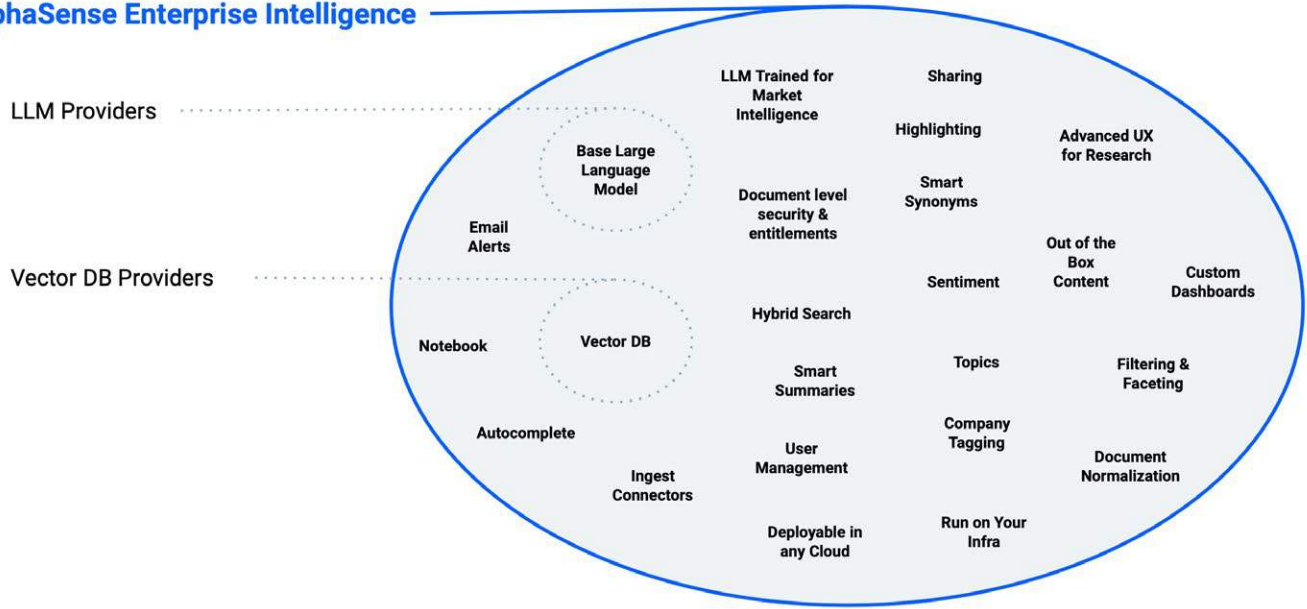
In fact, there's a variety of elements that need to be considered if you want to deliver an end-to-end solution to your organization.

Some of the key areas to consider are:

- Model training, including content sourcing for training purposes
- Document processing (ie., normalization, company name recognition, topic extraction, etc.), and permissioning capabilities
- Broader semantic search capabilities, relevancy algorithms, sentiment analysis and topic or theme or KPI identification—all of which require domain expertise to do well
- Scalability features including user management, broad system connectors, infrastructure compatibility, and collaboration or workflow tools
- Front-end interfaces and productivity tools cannot be underestimated when considering the broader change management that may be required

All of these play fundamental roles in how usable and scalable system solutions will be, and require distinct consideration and scoping. Ultimately, there's a lot more to consider than just what LLM to choose when evaluating how to implement genAI and improve the value of your internal content.

AlphaSense Enterprise Intelligence



AlphaSense's [Enterprise Intelligence](#) was developed by 500-plus developers in our product organization, focused on making internal content searchable. By leveraging 10+ years of AI and search expertise and investing millions of dollars into deciphering a solution to this challenge, we've developed a vertical LLM fine-tuned on business and financial content with a vector database.

Ultimately, buying leading LLMs or Vector Databases still require you to build, which is an extremely costly process that ranges from procuring external content, creating training data, and then training the LLM, as well as developing UI and feature sets. AlphaSense's Enterprise Intelligence is an out of the box solution that is ready now and adds value to your organization immediately.

AlphaSense is Your AI-Powered Insight Hub

AlphaSense is committed to unlocking the value of your internal content. Today, we help over 1,300 clients of ours to integrate their internal content alongside premium external business content—which includes 10,000 private, public, premium, and proprietary trust business sources, including leading equity research from 1,000+ sell-side and independent firms as well as 30,000-plus expert-led interviews—available within our platform. Customers can go into AlphaSense and, within our environment that boasts secure end-to-end encryption of client data, experience firsthand how our system processes documents for:

- Company name recognition
- Topic identification and extraction
- Sentiment analysis
- Synonym recognition
- Tables and charts

When you run a search on a related topic or company, your content will surface alongside relevant premium external content like broker research, industry reports, company documents, expert call transcripts, and more in the search results.

Further, we are building on these existing Enterprise capabilities with the announcement of our [Enterprise Intelligence solution](#), which allows you to search, discover, and interrogate your proprietary internal content alongside an out-of-the-box content library.

At AlphaSense, we have built a platform that is uniquely positioned due to our expansive universe of content. We boast

broker research, proprietary insights in our Expert Transcript library, and in our Enterprise Intelligence solution, help customers raise the value and utility of their internal content.

To bring that value into focus, lets run through three real-world value propositions:

Removing Intelligence Silos

Content can become siloed, making it difficult to harness firm-wide intellectual property. This silo effect is most often the result of firms who are acquisitive or combining/ reorganizing teams within a company, but have not yet fully integrated all of the research and IP that sits within these assets.

In some cases, the "silo effect" is intentional. An investment team, department, or division is intended to function independently and must never share information with other internal teams. In these cases, the User Management and fine grained access controls at the document level are critical.

Restricting access to certain documents is common to any firm and must be supported within any enterprise intelligence solution. More importantly, the concept of document entitlement must flow through to genAI summarization and Chat experiences, or else the user has a back door into content sets they are otherwise not permitted to access. This is a reality the AlphaSense team has often seen get overlooked by vendors.

Purpose-Built AI

For over 10 years, we've designed our artificial intelligence technology with a singular focus: to create solutions that deliver relevant, real-time insights for our customers.

We started with delivering AI-powered semantic search solutions like Company Recognition, Sentiment, Topic and KPI Extraction, and Smart Synonyms™ to help users find differentiated insights faster. Now we have GenAI-powered Smart Summaries, automatically creating SWOT analyses and summarizing documents. So, what is next?



As we expanded into genAI, our team has taken a measured approach to building our technology. Building off of our multi-layer proprietary AI stack that we have fine tuned over the past decade plus, and building off of foundational models, we have trained our large language model on an unparalleled set of business-focused content, using those guardrails to help focus the genAI learning.

The result is our proprietary, vertically-trained model; that delivers highly accurate, and relevant insights and KPIs. Out of the box, we speed up research workflows with instant summarization on earnings calls, company outlooks, SWOT perspectives, expert calls, and more. Features like deep citation links ensure users can easily verify the content delivered to them with a single click.

And when clients layer these features on their proprietary knowledge base, they start creating more tailored results and summarizations that seamlessly blend internal research as well as relevant external perspectives.

Enterprise-Grade Data Protection

When you explore layering technology on your valuable internal content and IP, security needs to be front and center.

To that end, AlphaSense is committed to the security and privacy of our clients' data. This includes offering end-to-end encryption of customer information, a zero-trust security model, flexible deployments, and robust entitlement awareness.


Our secure cloud environment complies with global security standards—SOC2, ISO270001 compliant, conducts regular, accredited third-party penetration testing. Additionally, we offer FIPS 140-2 standard encryption on all content, and support SAML 2.0 integration to promote user authentication. Specific to genAI, all queries and data fed into and generated by smart summarizations remain within the platform, with no third parties involved.

In addition to being available as an AlphaSense-managed solution in our secure cloud, Enterprise Intelligence can also be deployed via a customer-managed private cloud. This may appeal to customers who need to comply with more stringent internal information security policies or controls.

Integrations

AlphaSense provides easy integration with third-party connectors such as OneNote, Sharepoint, Box, and Evernote to enhance your productivity and streamline your collaboration. Not only can you upload files but manage tens of millions of documents by uploading directly through our API, 3rd party connectors, or customer-managed solution in the cloud of your choice. Our integrations also allow you to manage users and entitlements, and embed AlphaSense insights into other 3rd party products at scale.

Integrations are key in improving your team's workflow in a secure and automated way, and AlphaSense integrations are powered by AI and help capture what others easily miss. With integrations, you can easily search across all internal and external content to find actionable insights, as well as collaborate with team members to share those insights internally.



Benefits of AI Generally for Internal Content

Effective business decisions rely on efficient access to the right information. Organizations equipped with speed to insight surge ahead, surpassing those without this invaluable advantage. At AlphaSense, our commitment has been centered on quickly delivering critical information to our customers through the application of our exclusive AI technology within our platform.

We've meticulously developed and enhanced our core AI algorithm to attain a nuanced comprehension of business and financial language, which has resulted in the creation of a potent search engine tailored for the world's leading companies and investment firms. Further, we've amassed an extensive and dependable collection of business content to ensure that our clients are utilizing the world's most robust market intelligence search platform.

AI manifests within AlphaSense in the shape of tools tailored for today's business professionals to identify market opportunities and take the lead in in their respective industries:

Smart Synonyms™: This innovative technology unique to AlphaSense analyzes speech patterns across tens of millions of search documents, resulting in a robust library of synonymous words. The platform ensures you get all the relevant search results—without any of the excess noise or extraneous effort.

Theme Extraction: This tool allows users to extract the most important topics from earnings transcripts in the context of key metrics, such as QoQ increase in mentions, positive/negative sentiment, and overall mentions.

Trending Topics: Our Trending Topics tool is key for unlocking critical insights from a vast body of data—in seconds. This AI tool, accessible via the main dashboard, allows you to see the top 100 topics that are trending for a single company or a group of companies in a watchlist or industry.

Sentiment Analysis: Our technology utilizes [Natural Language Processing \(NLP\)](#) to uncover market perceptions about any given topic. Further, it color-codes positive and negative sentiment for easy recognition as you browse search results. These AI features are executed in seconds, saving analysts precious hours that would have previously been spent combing through multiple earnings call transcripts.

With this tool, users save time while identifying the most relevant information at scale across every single competitor's transcript. And because our model is trained specifically on financial language, it can pinpoint the most valuable and important company insights.

Table Explorer: For a number of professionals, a surplus of time is wasted on manually aggregating or calculating historical financials and metrics within Excel. Table Explorer eliminates the need to manually collect data so you can more quickly double-check the numbers in Excel. With AlphaSense, you can easily identify quantitative insights faster and expedite extracting valuable data from tables in SEC filings, [broker reports](#), and other content sets to help you build better financial models.

Company Name Recognition (CNR): CNR is our home-grown solution for recognition, inter-company disambiguation, and salience classification of company mentions across AlphaSense content. CNR paired with relevance scoring, Smart Synonyms™, and proximity search leads our proprietary AI to understand what or who you're looking for in any given business context.

How GenAI Further Improves Upon Internal Content

GenAI is already transforming market intelligence, as data once buried within static documents can be surfaced and synthesized in real-time, with just a few keystrokes. And at AlphaSense, we recognize how this technology can not only benefit your research, but how you search for the insights needed to fuel it.

Critical inputs into market intelligence research—for example, a firm's internal research, investment memos, client deliverables, strategy presentations, and meeting summaries—are often fragmented and inaccessible, resulting in lost opportunities and doubled work.

[AlphaSense's Enterprise Intelligence](#) is the solution to this problem, serving as a secure, end-to-end market intelligence solution that leverages genAI to unlock insights from premium external and proprietary internal content—in a single platform.

With Enterprise Intelligence, you'll be able to apply our proprietary AI stack over your internal content at scale, effortlessly by searching for content and quickly gleaning their context through our summarization feature. Follow up with questions like "do we have any research on Google's AI investment?", and receive back instant insights unique to your organization's perspective. Or expand the reach of these insights even further to include valuable external perspectives (ex. broker research, industry reports, expert insights, etc.) to create a 360° view of any topic.

By ingesting your internal knowledge base, our AlphaSense Large Language Model (ASLLM) can securely understand your organization's data and thus provide a comprehensive, unique, and ever-evolving view of your business and markets—while you and your team contribute the uniquely human attributes of context and discernment to the research process—together, producing more knowledge than ever before.

AlphaSense

About AlphaSense

AlphaSense is a market intelligence and search platform used by the world's leading companies and financial institutions. Since 2011, our AI-based technology has helped professionals make smarter business decisions by delivering insights from an extensive universe of public and private content—including company filings, event transcripts, expert calls, news, trade journals, and equity research. Our platform is trusted by over 3,500 enterprise customers, including a majority of the S&P 500. Headquartered in New York City, AlphaSense employs over 1,000 people across offices in the U.S., U.K., Germany, Finland, and India.

LEARN MORE AT ALPHA-SENSE.COM

