



SECRETS OF GEN AI SUCCESS

How leaders like Bayer and Siemens Energy use generative AI
to increase revenue, improve productivity and better serve customers

TABLE OF CONTENTS

3	The AI reality check
4	Strides toward success
7	Gen AI in action
8	TS Imagine adopts AI at scale to save 30% in costs and 4,000+ hours of effort
10	Siemens Energy's AI chatbot democratizes access to 700k+ pages of knowledge
11	Bayer's business teams make better, faster decisions with self-serve insights
13	Penske fuels associate retention and safety with better performance insights
15	Sigma increases sales efficiency with prescriptive recommendations
16	Compare Club turns untapped call transcripts into new ways to delight and engage members
17	CARTO simplifies wildfire risk analysis by explaining geospatial data in natural language
18	Terakeet helps customers pinpoint new market opportunities 98% faster
19	Coda's conversation-based AI platform allows customers to turn data into action
20	A Premier Venture Capital Firm turns third-party data into investment insights
21	IntelyCare prevents job postings from falling through the cracks
22	Johnnie-O improves accuracy of geocoding address data to better serve customers
23	Alberta Health Services ER doctors automate note-taking to treat 15% more patients
24	Advisor360 constantly gauges customer satisfaction using an automated pipeline
25	Your path to gen AI success

THE AI REALITY CHECK

When generative AI catapulted to fame in late 2022, hopes were limitless. It was going to supercharge teams by automating drudgery. It was going to unlock new business models, unearth revenue streams and lead to unfettered productivity. It was going to cure diseases and save lives. Gen AI was going to change *everything*.

Fast-forward to today, and gen AI has changed many things — but some of its biggest promises are still more aspiration than reality. While nearly **two-thirds of C-suite executives** say trust in AI drives revenue, competitiveness and customer success, according to a Salesforce survey, Harvard Business Review reports that **up to 80% of AI projects fail** to make it into production.

So why the disconnect? When we examine the obstacles organizations face in bringing their gen AI use cases into production, it's not hard to understand why many are falling short.

Roadblocks include:

- The cost and complexity of managing multiple AI platforms — and then moving and duplicating data across all these infrastructures.
- The operational burden of democratizing access to AI, equipping the right people with access to the right tools and the right training.
- The governance requirements to ensure that data and large language models (LLMs) are meeting heightened security standards.

Compound these challenges with the ever-changing AI toolset, which seems to grow daily, and AI efforts may start to feel more like futile attempts.

But there are still plenty of reasons to be bullish about AI.

Organizations across industries are finding ways to overcome these hindrances and unleash the potential of AI across their business. These leaders are lowering the AI barrier to entry by building on an already robust data foundation that prioritizes flexibility and helps them stay agile to meet evolving needs. They're managing costs, reducing operational burden and ensuring their data stays governed.

And they're achieving these goals and more with Snowflake's versatile, powerful platform that makes AI easy, efficient and trusted.

80% of AI projects fail to make it into production

STRIDES TOWARD SUCCESS

Organizations of all sizes and industries are turning to Snowflake's unified AI and data platform to make their generative AI goals a reality, whether it's increasing sales efficiency and accelerating sales cycles (**Sigma**) or democratizing access to insights for nontechnical teams as varied as finance, sales and demand planning (**Bayer**).

These premier enterprises are succeeding by tapping into the power of **Snowflake Cortex AI**, a fully managed service that hosts and serves top-tier LLMs, allowing businesses to easily and efficiently develop generative AI applications without moving data outside of Snowflake's trusted boundary. In doing so, these organizations are building AI applications that are capturing tangible value — from pinpointing new market opportunities for customers 98% faster (**Terakeet**), to achieving 30% cost savings through greater efficiency (**TS Imagine**).

TERAKEET

98%

Faster in pinpointing new
market opportunities

TSIMAGINE

30%

Cost savings through
greater efficiency

“

A lot of organizations right now are struggling with whether generative AI is worth the expense and risk. Snowflake Cortex AI provides a solution that doesn't require initial investment or risk, and that's what has made our gen AI projects successful.”

—VISHWA RAM

Vice President, Data Science and Analytics,
Penske Logistics



THREE KEY USE CASES MAPPING THE WAY

While organizations are exploring a variety of new gen AI use cases, there are three primary areas where customers are making big strides with gen AI:

1. Text processing at scale:

Organizations are using Snowflake Cortex AI to gain access to industry-leading, open source LLMs, such as Mistral and **Snowflake Arctic**, that quickly perform complex tasks like translation, summarization and sentiment analysis. This processing takes place across massive volumes of text — from summarizing hundreds of pages of documents; to calculating sentiment across thousands of emails, calls and support tickets (**Advisor360**); to reformatting jumbled, unintelligible raw data into something clean and usable (**Johnnie-O**, **IntelyCare**) — all without leaving Snowflake's secure, governed perimeter.

A common way organizations are generating value with this use case is by empowering sales teams. **Sigma** uses Cortex AI to turn hundreds of Salesforce opportunities, support calls and chats into invaluable reports for sales that predict a customer's propensity to purchase, summarize win/loss insights and suggest next best actions. And companies, such as **Penske** and **Compare Club**, are also using these capabilities to breathe newfound value into untapped assets like customer reviews, support tickets, call transcripts and supplier contracts. **Alberta Health Services** is even using Cortex to improve efficiency in its emergency departments by automatically summarizing physician-patient interactions using an LLM, enabling doctors to see 10%-15% more patients per hour.

2. Conversational assistant:

Thanks to conversational assistants, enterprises are allowing employees to search documents and multimedia through chat. This capability gives teams intuitive, interactive access to information and insights while eliminating the need to dig through files or multiple applications.

Organizations like **Siemens Energy** are using Cortex AI to build these powerful chatbots based on a retrieval augmented generation (RAG) architecture that enhances data access and improves employee productivity, while tech platform **Coda** uses Cortex AI and its integrated features like Cortex Search (in public preview) to provide hyperlinked citations that verify the accuracy of its search results.

3. Automated business intelligence (BI):

Automated BI enables employees to perform ad hoc data analytics without writing SQL, which is typically used for generating analytics and interacting with databases. With powerful Cortex AI features like Cortex Analyst (in public preview) and Snowflake Copilot, users leverage intuitive, conversational interfaces that streamline data exploration, visualization and reporting to make data insights available to everyone — not just technical teams.

Bayer, for example, allows business teams to self-serve insights by using Cortex Analyst to convert natural language to SQL. This helps teams quickly and easily find valuable information around product performance, profits and beyond to make more strategic decisions. Meanwhile, **Advisor360** uses an automated pipeline to gauge and track customer sentiment, saving at least a day's worth of work every month for a senior data scientist — and at a fraction of the cost of an alternative solution.



Accelerates sales cycles

and delivers a better customer experience



Enhances strategic decision-making

with self-serve data insights

THE JOURNEY PAYS OFF

These real-world gen AI use cases are delivering real results for companies both big and small. In doing so, these organizations are capturing tangible value from their AI use cases — from **Siemens Energy** accelerating research and development and time to value, to **TS Imagine** saving over 4,000 hours of manual effort a year while cutting costs by nearly a third.

And the enterprise innovation extends far beyond these pages. Here are just a few other organizations using Cortex AI to achieve their business goals:

- A leading global video hosting platform uses LLMs in Snowflake Cortex AI to generate new insights and sentiment analysis from video transcripts, helping significantly boost conversion rates by better tailoring communications to customers.

- A multibillion-dollar American equipment rental enterprise uses the Mistral LLM in Cortex AI to cost-effectively analyze thousands of sales notes and reduce customer churn.
- A Fortune 50 insurance giant uses Cortex AI to accelerate recruitment efforts by adding more structure to resumes, in addition to building an employee chatbot that quickly surfaces information sprinkled across hundreds of manuals.

As these businesses illustrate, AI has enormous potential that can extend into every team in the enterprise. While the evolving landscape will continue to spur new obstacles and opportunities, building on Snowflake's robust AI platform can help your organization stay one step ahead to realize the promise of AI.

SIEMENS
energy

**Eliminates years'
worth of manual
effort for
faster time
to value**

TSIMAGINE

**Saves over
4,000 hours
a year on tedious
processes**

GEN AI IN ACTION

Future-minded leaders across industries are building easy, efficient and trusted gen AI applications with Snowflake.



See how the organizations in the following pages are using the AI Data Cloud and Cortex AI to achieve their goals — from better serving customers to gaining competitive advantage.

TS IMAGINE ADOPTS AI AT SCALE TO SAVE 30% IN COSTS AND 4,000+ HOURS OF EFFORT

By automating email monitoring and streamlining management for support and data tickets, TS Imagine reduces manual tasks, optimizes data processing and improves efficiency — saving time and cutting costs.

TSIMAGINE

INDUSTRY Financial technology

USE CASE Text processing; conversational assistant

TS Imagine delivers a SaaS platform for integrated electronic front-office trading, portfolio management and financial risk management. Traders, risk managers and operations executives at more than 500 global financial institutions use TS Imagine products so they can better focus on generating and protecting assets within today's fast-evolving markets.

CHALLENGE

Before Snowflake, TS Imagine's homegrown data solutions required redundant, manual work. Every year, members of the data management team spent thousands of hours reading more than 100,000 emails to monitor for notifications from their data providers. These critical emails detailed upcoming product changes, any of which could ultimately impact TS Imagine's clients. Missing even a single email could lead to a downstream product outage — and very unhappy customers.

Another pain point was prioritizing the more than 60,000 support and data tickets that the customer team received annually. Since TS Imagine comprises two legacy firms founded in the 1980s, finding internal information about the combined company post-merger was also inefficient and time-consuming for current employees and new hires, with knowledge spread out in far too many places. Finally, employees had to peruse dense terms and conditions documents published by derivatives exchanges about thousands of financial instruments to produce structured data for TS Imagine's trading and risk management products.

IMPACT

- **30% cost savings** — without sacrificing performance — by using Snowflake vs. other industry-leading, pretrained LLM APIs
- **Ease of use for data teams** to go from traditional natural language processing to generative AI at scale in only 6 months
- **Massive time savings** for employees, from over 4,000 hours spent on manual email monitoring tasks to near-instant, AI-powered categorization for just one use case

[Read the full story](#)



RESULT

TS Imagine has operationalized AI at scale with Snowflake Cortex AI. Instead of employees spending over 4,000 hours a year on an error-prone sorting process, TS Imagine has automated email intake, using Cortex AI to delete duplicate or nonrelevant messages, and create, assign, prioritize and schedule JIRA tickets for each email. The team hasn't missed a notification since implementing the new process in December 2023. And those 60,000+ customer support and data tickets? Cortex AI classifies them by sentiment, complexity, urgency, clarity, uniqueness and impact so client managers and support team members can resolve customer-facing issues much faster.

TS Imagine uses Cortex AI to improve internal efficiencies, too. In-depth knowledge about TS Imagine's products — found in tens of thousands of articles, documents and support tickets within the company's knowledge base — is now all available via LLMs in Cortex AI, including Snowflake Arctic, Llama and Mistral. Employees find info about the company fast by simply asking questions of a chatbot, powered by [Streamlit in Snowflake](#). And rather than painstakingly reading terms and conditions within fact sheets about financial instruments, teams rely on Snowflake Arctic to transform such PDFs into easily digestible structured data, which then automatically feeds into TS Imagine's products.



4,000+ HOURS SAVED



30% COST SAVINGS



With Snowflake, I can empower smart people to bring AI to life in one place, even if they're not tech experts. Cortex AI is a one-stop shop. It scales, it's easy, and the data stays 100% in Snowflake's secure environment."

—THOMAS BODENSKI

COO & Chief Data & Analytics Officer,
TS Imagine

SIEMENS ENERGY'S AI CHATBOT DEMOCRATIZES ACCESS TO 700K+ PAGES OF KNOWLEDGE

With Snowflake Cortex AI and Streamlit in Snowflake, Siemens Energy built an AI chatbot to quickly surface and summarize over half a million pages of internal documents, helping the R&D team maximize productivity and effectiveness while accelerating decision-making and competitiveness.



INDUSTRY Manufacturing

USE CASE Conversational assistant

A global leader in the energy business, Siemens Energy's innovative solutions contribute to around one-sixth of the world's energy production. Half of Siemens Energy's portfolio is based on technology pertinent to a decarbonized energy world, demonstrating the company's deep commitment to sustainable energy technology.

CHALLENGE

Siemens Energy has a deep knowledge base contained in more than 700,000 pages of proprietary research and development documents. The company wanted to make the valuable information in these paper documents available to its globally distributed research and development (R&D) team — but reading all this information would take someone a staggering four years. The challenge was to not only digitize the information, but to then make it more accessible to teams so they could use the wealth of knowledge to accelerate discovery, collaborate more effectively, boost decision-making and fuel innovation.

SOLUTION

Applying a comprehensive citizen development approach, Siemens Energy built an intuitive document Q&A chatbot in Snowflake that democratized knowledge discovery and collaboration in a secure, highly economical way. Developed in a very short period of time, this state-of-the-art chatbot uses a retrieval augmented generation (RAG)-based architecture — built using Streamlit in Snowflake and Snowflake Cortex AI — that made the internal company documents securely available to off-the-shelf large language models without the need for training, fine tuning or moving data outside Snowflake's governed platform.

Whereas they had to previously spend hours sifting through and reading research papers, researchers now simply use the chatbot to quickly identify relevant research excerpts with accompanying hyperlinks to the source material. This breakthrough helps Siemens Energy save years' worth of manual searching while ultimately speeding time to market and sharpening the company's competitive edge.

IMPACT

- Information from 700K+ pages of dense documents summarized and available to teams, which **accelerates R&D, innovation and time to market**
- ~25 highly specialized R&D engineers now have **near-instant access to data and insights**
- Hours of time saved — amounting to **years' worth of effort** — by eliminating the need to manually search for data

[Watch their presentation](#)

“ This chatbot, built with Snowflake Cortex and in combination with our comprehensive citizen development approach, has been a game changer for Siemens Energy, particularly in the realm of fast and economic data democratization. It has optimized our resource allocation, saving many manual hours spent looking for information. The ability to unlock and democratize the data hidden in our data treasure trove has given us a distinct competitive edge. It's intuitive, it's innovative, and it's driving us forward in our mission — including achieving high scalability for further use cases already in progress.”

—TIM KESSLER
Head of Data, Models & Analytics, Siemens Energy

BAYER'S BUSINESS TEAMS MAKE BETTER, FASTER DECISIONS WITH SELF-SERVE INSIGHTS

Bayer's Consumer Health North America team converts natural language queries into accurate and actionable SQL, helping nontechnical users — from customer finance to sales to demand planning — quickly analyze data for faster decision-making, without relying on data teams.



INDUSTRY Healthcare and life sciences

USE CASE Automated BI

Guided by the mission “Health for all, Hunger for none,” Bayer is a global life sciences enterprise that designs products and services to help tackle some of the world’s biggest challenges while serving the most essential human needs of health and nutrition.

CHALLENGE

Businesses like Bayer move fast, especially when partnering with third-party retailers and suppliers as part of their go-to-market strategy. To make quick, informed decisions about the business, teams across the company — from customer finance to sales to demand planning — need near real-time, accurate insights from their data. Yet most teams only had access to this data via dashboards, which often resulted in information overload for users and cumbersome, time-consuming development for data teams. With limited coverage and flexibility, these dashboards forced users to manually dig through and interpret data, which hindered teams from quickly extracting insights, increased analytics inconsistencies and introduced the potential for human error.

IMPACT

- **More strategic decision-making** by giving business teams the ability to ask questions of Bayer’s curated data sets in natural language
- **Insights at scale** for users across the enterprise — from sales to customer finance to demand planning
- **Better productivity** for both technical and nontechnical departments by reducing reliance on data teams and enabling self-serve data access

[Watch their presentation](#)



SOLUTION

Leveraging Cortex Analyst to abstract database interactions by converting natural language to SQL, Bayer now empowers its business users to easily self-serve insights through an intuitive Streamlit in Snowflake chatbot interface. Instead of having to comb through dashboards or request support from data and analytics teams, departments like customer finance, sales and demand planning use natural language to seamlessly interact with analytical data stored in Snowflake tables. This helps them find answers faster while unearthing richer, actionable insights in easy-to-understand terms and visuals. For example, a customer finance user can request a tabular report of customer profit and loss statements by brand and month, or ask which products have the best and worst performance by net sales at a given retailer. Using Cortex Analyst's text2SQL capabilities, the chatbot returns accurate, reliable results in near real time to help teams better understand the business and make strategic decisions — no technical expertise or IT support required.

“The core of this capability is high-quality responses to a natural language query on structured data, used in an operationally sustainable way. This is exactly what Snowflake Cortex Analyst enables for us. What I'm most excited about is we're just getting started, and we're looking forward to unlocking more value with Snowflake Cortex AI.”

—MUKESH DUBEY

Product Management and Architecture Lead,
Bayer U.S.

PENSKE FUELS ASSOCIATE RETENTION AND SAFETY WITH BETTER PERFORMANCE INSIGHTS

Using Snowflake Cortex AI, Penske Logistics equips its managers with a detailed summary of talking points based on performance indicators obtained from a multitude of devices and operational data sources. The summaries enable a guided conversation with associates (employees) in a natural language format for more effective, informed meetings that help increase associates' efficiency, safety and retention.



INDUSTRY Logistics and supply chain management

USE CASE Batch text processing

Penske is a leader in the fields of supply chain and logistics, providing dedicated transportation, warehousing, transportation management, freight brokerage and related customized supply chain solutions to its customers in North America and around the globe. Penske Logistics currently operates a vast logistics network with a workforce of more than 20,000 associates worldwide, and it is among the top-ranked service providers of its kind.

CHALLENGE

Safety, performance and retention of associates are paramount for Penske Logistics. Managers aim to improve these areas with personalized insights and recommendations to their teams informed by massive amounts of data. However, insights from this data were not easily accessible and analysis was a manual process to parse through nearly 30 data columns in a dashboard.

To help provide managers with a comprehensive view of performance indicators, the data science team explored various AI solutions that could analyze and consolidate data to look for trends and proactively raise awareness to spark conversation.

IMPACT

- Greater safety, performance and retention thanks to clearer, more precise insights and action items
- Better productivity for managers across weekly and monthly performance conversations
- Thousands of columns of data covering two product lines automatically synthesized and transformed into digestible talking points

[Read the full story](#)



SOLUTION

Penske Logistics found traction with Snowflake's gen AI. Because it hosts and serves large language models (LLMs) within Snowflake's governed perimeter, Cortex AI made it easy and cost-effective for the team to securely bring LLMs to their machine learning modeling results — without requiring additional approvals or investment for new AI vendors. Now, Penske uses LLMs in Cortex AI to process performance indicators in record time and produce a plain-English summary of the key talking points and recommendations for each associate. These insights allow managers to guide conversations with more clarity, precision and actionable next steps. Since rolling out the solution, managers are enjoying productivity gains from reading a comprehensive paragraph, eliminating multiple sources and columns of data. The company is already seeing improvements in associate performance, safety and retention.



A lot of organizations right now are struggling with whether generative AI is worth the expense and risk. Snowflake Cortex AI provides a solution that doesn't require initial investment or risk, and that's what has made our gen AI projects successful."

—VISHWA RAM

Vice President, Data Science and Analytics,
Penske Logistics

SIGMA INCREASES SALES EFFICIENCY WITH PRESCRIPTIVE RECOMMENDATIONS

Armed with analyses like an account's propensity to purchase and "why we win" summaries, go-to-market teams save time, accelerate sales cycles and deliver a more tailored experience to prospects and customers.



INDUSTRY Technology

USE CASE Text processing and summarization

Sigma's powerful business intelligence platform scales billions of records using spreadsheets, SQL, Python and AI — without compromising speed and security. Its solution helps users analyze information via dashboards, then gain access to the underlying data in an intuitive, spreadsheet-like interface. This interface acts as a frontend to cloud data platforms like Snowflake and provides the average spreadsheet user with a simple way to interact with LLMs in Cortex AI.

CHALLENGE

Sigma helps teams of all sizes, from startups to Fortune 500s, better understand and interact with their data. But to deliver on this goal — and continue helping more companies — Sigma's own internal teams require a deep understanding of both their customers and prospects. Sales teams needed a scalable way to easily access data while pinpointing the most promising leads so they could prioritize efforts and maximize their time. At the same time, Sigma's go-to-market and operations teams needed a more reliable way to stitch together account details rather than relying on individuals' notes, which could be incomplete, inconsistent or prone to human error.

SOLUTION

Sigma now uses Cortex AI to fuel productivity and insights for teams across the business. Through Cortex AI, Sigma taps into the power of LLMs like Mistral and Snowflake Arctic to predict a prospect's propensity to purchase and equip sales teams with suggested next-best actions to accelerate the sales cycle. Business teams across Sigma can also ask questions of the company's data to better understand customer needs on a higher level, such as "why we win" summaries and customer sentiment analysis from hundreds of Salesforce opportunities and support calls and chats. Available within an internal Sigma app, these insights help sales and operations teams focus their efforts, provide better service and understand customer needs on a deeper level — benefits that have made this innovation wildly popular among Sigma's teams.

IMPACT

- **Reduced time and effort for sales teams** to identify the most promising prospects by using AI to predict an account's propensity to purchase
- **Better understanding of sales motions** and successes by equipping teams with win/loss summaries
- **Access to rich, easy-to-digest insights** around customer preferences, needs and sentiment by summarizing hundreds of Salesforce opportunities, account team notes and support calls and transcripts

[Watch their presentation](#)

“Snowflake's Cortex AI has immediately unlocked new workflows for our team. When combined with Sigma's Input Tables in our data applications, we're turning these model outputs into real value faster than ever before.”

—JAKE HANNAN,
Head of Data, Sigma

COMPARE CLUB TURNS UNTAPPED CALL TRANSCRIPTS INTO NEW WAYS TO DELIGHT AND ENGAGE MEMBERS

Compare Club uses Cortex AI to efficiently summarize hundreds of thousands of pages of call transcripts, extracting valuable insights that improves member experiences and team productivity while keeping customer data secure.



INDUSTRY Financial services

USE CASE Text processing

Compare Club helps millions of Australian consumers make more informed purchasing decisions on products and services across health and life insurance, energy, home loans and more.

By treating data as a valuable tool, Compare Club experts make it easier and faster for customers to sign up for the most appropriate product and get more value from their budgets.

CHALLENGE

Providing an exceptional, personalized experience to customers is critical for Compare Club — especially for returning members, who are more likely to make a purchase. Customer calls are an important vehicle to deliver this experience, yet complex details from these conversations were not always recorded in the company's CRM, making it difficult to use this information in future calls. This oversight caused frustration and prevented Compare Club's experts from accessing the complete customer history they needed to provide the most effective recommendations. While call transcripts contained a trove of valuable insights, the team would have had to manually sift through thousands of transcripts — each exceeding 100 pages — which wasn't feasible from a time and bandwidth perspective.

Compare Club wanted to use generative AI to summarize these transcripts, but document length presented challenges to various LLMs' context window limits. And running a summarization task for every call meant having to move large volumes of sensitive call recordings into a third-party service, which would require validation and additional security approvals.

SOLUTION

Compare Club quickly overcame these challenges by using Cortex AI to run LLMs securely inside Snowflake, eliminating the need to move data while easily running both preprocessing and LLM tasks with a bit of SQL and Python. Now, Compare Club efficiently equips business teams with valuable insights extracted from hundreds of thousands of transcript pages, including details like customer goals, needs, objections, loyalty, history and enthusiasm. These nuances help Compare Club teams — from sales to support to customer success — better serve and engage repeat members to improve their experience and retention.

IMPACT

- **Enhanced member experiences** thanks to fresh insights from thousands of lengthy call transcripts
- **Greater productivity and effectiveness** for teams like sales and support with easily accessible, comprehensive customer details
- **Peace of mind for sensitive information** by keeping customer data within Snowflake's secure, governed perimeter

[Learn how they built it](#)



CARTO SIMPLIFIES WILDFIRE RISK ANALYSIS BY EXPLAINING GEOSPATIAL DATA IN NATURAL LANGUAGE

By turning complex geospatial data into clear, natural language explanations and graphics, CARTO makes it easier for nontechnical users to understand complex predictions like wildfire risk scores.

CARTO

INDUSTRY Technology

USE CASE Text processing and summarization

As the leading location intelligence platform, CARTO aims to make geospatial analytics accessible to all audiences. By presenting data visually, the company helps organizations in industries like urban planning, environmental management and business analytics drive better outcomes — from more efficient delivery routes to better behavioral marketing.

CHALLENGE

CARTO helps users create geospatial map dashboards, models and workflows, which are essential for visualizing complex spatial data and uncovering geographic patterns that inform strategic decision-making. But when a data scientist creates and maps predictions from a particular machine learning model, nontechnical users often find it challenging to understand the rationale. For example, if a homeowner receives a letter informing them that insurance premiums are increasing due to a fire risk index, they want to understand the cause. This also means that insurance companies are under mounting pressure to provide transparency and clear explanations around how they're rating different locations.

SOLUTION

With Cortex AI, CARTO translates complex metrics used to decide which ML model to use into a more accessible, easier-to-understand explanation for a nontechnical user. Building on the wildfire example, CARTO created a composite index using geospatial data to rank the wildfire risk of a given area on a scale of “no risk” to “very high.” Using Cortex AI functions, CARTO then summarizes and explains this map data in natural language for users, such as showing that an average annual temperature increase of three degrees has elevated fire risk in their region. To make this data even more accessible, CARTO visualizes the results in an intuitive Streamlit in Snowflake dashboard that allows users to interact with the data.

IMPACT

- **More transparency** by translating complex geospatial data into natural language and interactive Streamlit in Snowflake dashboards
- **Better understanding for nontechnical users** seeking to understand complex data questions like how insurance premiums are calculated
- **Greater productivity for data scientists** by automating explanations that accompany their analysis

[Learn how they built it](#)



TERAKEET HELPS CUSTOMERS PINPOINT NEW MARKET OPPORTUNITIES 98% FASTER

Using advanced LLMs to assess market opportunities faster, Terakeet has slashed analytics time while enhancing security by eliminating the need to move sensitive data.



INDUSTRY Technology

USE CASE Text processing

Terakeet is the preferred Owned Asset Optimization (OAO) partner for Fortune 500 brands seeking to connect with customers and grow their business online. Terakeet helps brands optimize and unify their digital assets, such as company values, websites and market positioning, to meet consumers as they search for solutions.

CHALLENGE

Terakeet helps its customers make total addressable market (TAM) estimations to understand potential revenue opportunities for their products or services, and effectively target untapped markets or customer segments. These assessments help brands understand the total addressable market of keywords and topics most relevant to their target audience, so they can tailor their marketing and development strategies accordingly. To build these assessments, Terakeet starts from a client's online properties. After pulling the client's and competitors' search landscape, the full set of results contains keywords and topics that may not be relevant. Prior to the TAM project, refining these results required a variety of methods, including tokenization, term matching and topic modeling that used an older natural language processing model. But these methods were cumbersome and time consuming, requiring significant manual effort.

SOLUTION

The text processing capabilities of Snowflake Cortex AI let Terakeet modernize and accelerate the TAM assessment process. Cortex AI has prebuilt task-specific functions that use LLMs behind the scenes to perform complex tasks like translation and summarization. With easy access to embedding functions in Cortex AI, the data science team transforms textual data into numerical vectors — or embeddings — which can identify data patterns and trends like semantic similarity. The team now performs these calculations directly in the Snowflake platform, which has eliminated the need to move data to third-party services outside of Snowflake's secure perimeter. This has not only resulted in significant time savings, reducing the assessment process from 48 hours to 45 minutes, but it has also enhanced data security. With such marked results, Terakeet's marketing and sales operations teams now use Cortex AI's capabilities for all of their clients and potential customers.

IMPACT

- **98% faster TAM calculations** — from 48 hours to 45 minutes
- **Increased customer satisfaction** by identifying untapped market opportunities more efficiently
- **Enhanced security** by eliminating the need to move data and building gen AI on Snowflake's governed platform

“Snowflake Cortex AI has changed how we extract insights from our data at scale using the power of advanced LLMs. Our teams can now quickly and securely analyze massive data sets, unlocking strategic insights to better serve our clients. We've reduced our processing times by 64x with the power of Snowflake's new AI features.”

—JENNIFER BRUSSOW
Director of Data Science, Terakeet



CODA'S CONVERSATION-BASED AI PLATFORM ALLOWS CUSTOMERS TO TURN DATA INTO ACTION

Coda Brain helps customers interact with their data through a conversational assistant, providing easy, actionable access to both structured and unstructured data.



INDUSTRY Technology

USE CASE Conversational assistant

Coda's technology transforms how businesses interact with and leverage their data. Coda's unified platform blends the flexibility of documents, the structure of spreadsheets, the power of applications and the intelligence of AI into a single interface.

CHALLENGE

Coda's unified platform serves as a central hub of knowledge for customers. To empower companies with the ability to extract actionable insights from their vast amounts of data, Coda wanted to build conversational services on top of its platform. However, traditional enterprise tools are not equipped to handle the unique challenges posed by enterprise data. The services needed to navigate complex permission structures, deliver precise and verifiable answers, and adhere to stringent enterprise-level security and governance requirements.

SOLUTION

Built on Snowflake and using Snowflake Cortex AI, Coda Brain is a turnkey RAG platform that works as a conversational assistant for enterprise customers. Coda Brain helps users ask questions of their enterprise data — structured or unstructured — and get reliable results to turn data into action across the enterprise. Teams no longer need to spend countless hours searching for enterprise information, and Coda Brain uses Snowflake's robust role-based access control (RBAC) to ensure data stays secure. To minimize hallucinations and increase the accuracy of search results over text-based documentation, Coda uses Cortex Search and provides citations to the data contained in responses. And by using Cortex Analyst to turn natural language text into analytical answers via SQL and accessing structured data, Coda further democratizes access to insights.

IMPACT

- **Increased enterprise efficiency and productivity**
- **Verifiability of knowledge** with citations through Cortex Search
- **Enhanced security and governance** with Snowflake's granular data permissions and RBAC

[Watch their presentation](#)

“Snowflake Cortex AI forms all the core building blocks of constructing a scalable, secure AI system. Coda Brain uses almost every component in this stack: The Cortex Search engine that can vectorize and index unstructured and structured data. Cortex Analyst, which can magically turn natural language queries into SQL. The Cortex LLMs that do everything from interpreting queries to reformatting responses into human-readable responses. And, of course, the underlying Snowflake data platform, which can scale and securely handle the huge volumes of data being pulled into Coda Brain.”

—SHISHIR MEHROTRA
CEO, Coda

A PREMIER VENTURE CAPITAL FIRM TURNS THIRD-PARTY DATA INTO INVESTMENT INSIGHTS

This financial firm uses LLMs to securely analyze and summarize large volumes of data, providing investment insights without the need for additional infrastructure or resources.

INDUSTRY Financial services

USE CASE Text processing and generation

Focusing on software and internet investments, this premier American venture capital firm has been an early investor in several leading technology companies.

CHALLENGE

This leading firm wanted to explore leveraging AI and LLMs to improve efficiency and productivity for several functions in the business, including marketing, research, portfolio impact and human resources. But the firm's small technology team was made up of just two data scientists and two full-stack engineers. The company lacked the resources to spend time vetting new vendors or to access, orchestrate and deploy pipelines using LLMs. Given the confidential nature of the business, the team also had to ensure the highest levels of security and governance for its data and AI.

SOLUTION

Using its existing, production-ready stack of data build tool (dbt) and Snowflake, the data sciences team turned to Snowflake Cortex AI to create value for the business without onboarding a new vendor or spending money on additional infrastructure, resources or upskilling. With access to industry-leading, open source LLMs within Snowflake Cortex AI, the team now runs text analytics across millions of records to gain deeper insights into both public and private companies. The firm has streamlined sourcing for investment teams by better categorizing and summarizing company data, such as company type and job title classifications, as well as providing investment teams with LLM-generated recommendations for similar companies by measuring statistics like web traffic growth.

IMPACT

- Cutting-edge LLMs available — **zero new vendor onboarding required**
- **No dependency on engineers** to access, orchestrate and deploy pipelines using LLMs
- **Instantly available, security-vetted environments** to process proprietary marketing and sales database

“Snowflake Cortex AI has made our database more robust with AI-enabled categorizations and classifications. We've improved our due diligence, spun up competitive market dashboards more quickly, and enhanced similar company and talent searches. Our data scientists now use Cortex AI to get some pretty exciting results with minimal effort and time investment.”

—HEAD OF DATA PRODUCTS
Venture Capital Firm



INTELYCARE PREVENTS JOB POSTINGS FROM FALLING THROUGH THE CRACKS

Using Cortex AI, the healthcare staffing agency used LLMs to efficiently extract pertinent information from thousands of job listings, resulting in an enhanced experience and better matches for candidates.



INDUSTRY Healthcare and Technology

USE CASE Text processing

IntelyCare provides a comprehensive platform that helps match healthcare organizations with qualified nursing professionals for open positions, allowing patients to get the care they need and nurses to find the roles that better fit their lives.

CHALLENGE

With hundreds of thousands of open jobs at any given time, filling those roles begins with organizing the postings, which is much more complicated than it sounds. While IntelyCare has direct relationships with many organizations, many of the largest healthcare systems prefer to post openings through a vendor management system (VMS), which can be accessed only by vetted agencies like IntelyCare. To be able to include these opportunities in its database and app, however, IntelyCare must process each post in a way that is organized and relatively uniform. That, of course, presents a challenge, given that each VMS adheres to its own system of standard practices. It's not uncommon for fields to be left blank on some posts or for bodies of text to be unintelligible to IntelyCare's internal tools; as a result, more than 30% of job posts were getting lost in processing.

RESULT

IntelyCare began using LLMs in Cortex AI to quickly extract pertinent information, both simple and complex, from these thousands of posts: specialty, pay range, specialized years of experience required, whether local applicants can apply for travel positions and so on. Then IntelyCare can organize the posts in thoughtful ways and without fear of losing opportunities because of, say, incompatible formatting — pushing that 30% of lost job posts down to zero. Not only has that enhanced the user experience for applicants, but it's also lightened the burden on IntelyCare's recruiters, who would tirelessly sift through individual postings to find the best matches for their candidates.

Cortex AI also helps IntelyCare manage and process the constant stream of postings it fetches from public job boards. Instead of dealing with complex orchestration and frequent maintenance, or having to build and train new models just to add new tags to jobs for easy categorization on their pipeline, now the team can replace thousands of lines of Python code with a single Cortex query — all while delivering an improved customer experience.

IMPACT

- **30% more job postings processed**, thanks to Cortex AI's ability to extract and organize jumbled information
- **Enhanced user experience** for candidates and recruiters by more thoroughly organizing job information and descriptions
- **Streamlined complex pipelines** used to appropriately categorize public postings into a single Cortex AI query



JOHNNIE-O IMPROVES ACCURACY OF GEOCODING ADDRESS DATA TO BETTER SERVE CUSTOMERS

The East-Coast-prep-meets-West-Coast-casual ecommerce brand used Cortex AI to clean up jumbled address data to gain better geographical insights.



INDUSTRY Retail

USE CASE Text processing

Clothing brand Johnnie-O designs classic, quality apparel that blends East Coast tradition with casual West Coast style. With a robust ecommerce business and a strong presence in premium clothing retailers nationwide, Johnnie-O delivers timeless pieces to wear from the golf course to the cocktail bar and beyond.

CHALLENGE

Like many largely ecommerce businesses, Johnnie-O understands the value in a simple shipping address. Receiving that address information from a variety of sources — not only from direct sales on its website but from wholesalers and dropshippers — the company could not accurately geocode a significant number of them, preventing the team from accessing relevant customer data that could help target marketing efforts. Typically, the company runs raw address data through an application that delivers geographic coordinates, which then makes it easy to link to U.S. census data. But for Johnnie-O, a significant number of these addresses failed for one reason or another; something as small as a typo or information in the wrong field would lead the application to fail.

RESULT

Instead of manually cleaning up these hundreds of thousands of data points, which would take many hours, the company looked to Cortex AI to automatically reformat the messy address data. Feeding these incompatible addresses into Cortex AI, Johnnie-O found that the Llama LLM it used was able to decipher jumbled data accurately and return clean addresses with ease. The company has now incorporated Cortex AI directly into a customer segmentation pipeline, automatically picking up addresses that the geocoding application could not handle, cleaning them up and then sending them back to the app again. The solution has slashed that failure rate to just 2%.

IMPACT

- **Reduced the failure rate of geocoding address data to 2%,** so typos or incorrect formatting are less of an issue
- **Enhanced market segmentation efforts and insights** with a more accurate demographic picture of its consumer base
- **Saved many hours** that would've needed to be spent to manually clean up address data

“Cortex AI is so easy to use and implement, especially because all of our data is already in Snowflake. Snowflake and Cortex AI have become the center of everything for us.”

—RICARDO LOPEZ,
Analytics Engineer, Johnnie-O



ALBERTA HEALTH SERVICES ER DOCTORS AUTOMATE NOTE-TAKING TO TREAT 15% MORE PATIENTS

The integrated health system of Canada's third-largest province improved physician workload and patient experience by using LLMs to generate summaries of doctor-patient interactions.



INDUSTRY Healthcare

USE CASE Summarization

Serving a province of 4.5 million residents, Alberta Health Services, with its more than 100 hospitals and 11,000 practicing physicians, aims to deliver patient-focused, quality healthcare that is accessible and sustainable for all Albertans. Its emergency departments alone get nearly 2 million visits per year, amounting to more than 5,000 a day.

CHALLENGE

Every visit with an emergency department doctor needs to be documented carefully, not only for regulatory purposes but also to communicate with specialists that may be needed or general practitioners who may be called to follow up. But typing up notes from every patient visit is a time-consuming and often tedious process. It also necessarily constrains doctors' ability to see more cases, which can ultimately lead to longer wait times for patients or seemingly rushed interactions. Furthermore, when practitioners aren't able to get to their notes right away, they could spend hours backlogging, and it's easy for them to forget small details when they've seen several patients back-to-back before doing documentation.

RESULT

One Alberta ED physician, also a trained software engineer, sought a way to automate his note-taking tasks by recording visits and calling an LLM to generate a summary. Seeing the potential of this use case, Alberta Health Services turned to Cortex AI to develop and run an app all within Snowflake's secure, governed environment. The app — which consists of recording patient interactions, transcribing and then summarizing them — has enabled doctors to spend more time in front of patients and less time in front of a computer.

Currently in its proof-of-concept phase, it's being used by a handful of emergency department physicians, who are reporting a 10%-15% increase in the number of patients seen per hour. That can ultimately translate into less-crowded waiting rooms, relief from overwhelming amounts of paperwork for doctors, even better-quality notes and higher-quality patient care. It also generates lengthier and often more empathetic discharge summaries, which improve the patient experience.

IMPACT

- **~15% more patients seen** in Alberta's emergency departments
- **Less time spent on paperwork** for physicians, freeing them to focus more on patient care
- **Enhanced patient experience** by giving more empathetic and lengthier discharge summaries than doctors have bandwidth to produce

“ This is a transformational change. You can begin to recognize that this technology, built with Cortex AI, can even start driving care, including suggesting and creating orders. There's a very bright future for this type of tool to foundationally change medicine, and it already is in some areas.”

—JASON SCARLETT,
Executive Director of Enterprise Data
Engineering, Alberta Health Services



ADVISOR360 CONSTANTLY GAUGES CUSTOMER SATISFACTION USING AN AUTOMATED PIPELINE

The wealth management software company used Cortex's built-in sentiment analysis to assess customer reception of its products quickly and with ease.



INDUSTRY Financial Services

USE CASE Sentiment analysis, text processing

Supporting financial advisors and wealth managers in nearly all aspects of their jobs, Advisor360's integrated and configurable SaaS platform informs business decisions, improves client service and even drives innovation for firms.

CHALLENGE

For SaaS companies like Advisor360, feedback from its users is as essential as it is constant, but getting a handle on all the inbound emails, calls and support tickets has historically been a tedious and largely manual process. To address that, the Advisor360° analytics and insights team had built a sentiment model from scratch, using highly specialized, Python-heavy code that would extract data and push it out to a file, then incorporate it into a dashboard. But the model required constant maintenance and updating, and when senior executives asked for insights incorporating the most up-to-date data, they'd often have to wait the better part of a day to get answers.

RESULT

After migrating all of its historical feedback data to Snowflake, Advisor360° created an automated pipeline using Cortex AI to cover the end-to-end process of gauging customer sentiment. Rather than taking a matter of weeks to establish the pipeline as anticipated, it took the team just two days.

Now, processing feedback occurs automatically every morning, and Advisor360° estimates that it saves at least a day's worth of work per month for a senior data scientist. What's more, the company found that using Cortex AI's built-in sentiment analysis functionality was about 1/25th the cost of the alternative, which would require using specialized virtual machines and calling in a larger and pricier LLM to do essentially the same thing. Advisor360° now has instant access to the most up-to-date customer insights — allowing the firm to provide the enterprise-class customer care it is known for.

IMPACT

- **Automated pipelines to gauge customer satisfaction** using built-in sentiment analysis capabilities in Cortex AI
- **Saved a day's worth of work per month** that a senior data scientist would otherwise spend maintaining complex pipelines
- **Achieved significant cost-savings** by using Snowflake's built-in tools, which came in at 1/25th the price of alternative solutions that required specialized virtual machines and more expensive LLMs

“When our data scientists tried to create this pipeline using native Azure tools, it probably took a solid month to figure out how to get permissions right and get through our firewall. But Cortex AI worked out of the box, integrating into our system seamlessly and translating into huge productivity gains for the team.”

—MARK COLEMAN,
Director of Analytics and Advanced
Insights, Advisor360

YOUR PATH TO GEN AI SUCCESS

The customer stories in this book are just the beginning of what gen AI can do for organizations using Snowflake's easy, efficient and trusted AI data platform.

Whether you're just starting to experiment with generative AI or already have your first AI applications in production, Snowflake enables you to extend the benefits of AI to your entire organization by streamlining development, reducing complexity and helping secure your data every step of the way.

Discover how Snowflake can help your organization be more innovative and competitive through generative AI. Visit Snowflake's [Gen AI webpage](#) for more information and hands-on resources to help you get started.





ABOUT SNOWFLAKE

Snowflake makes enterprise AI easy, efficient and trusted. Thousands of companies around the globe, including hundreds of the world's largest, use Snowflake's AI Data Cloud to share data, build applications, and power their business with AI. The era of enterprise AI is here.

Learn more at snowflake.com (NYSE: SNOW)



© 2025 Snowflake Inc. All rights reserved. Snowflake, the Snowflake logo, and all other Snowflake product, feature and service names mentioned herein are registered trademarks or trademarks of Snowflake Inc. in the United States and other countries. All other brand names or logos mentioned or used herein are for identification purposes only and may be the trademarks of their respective holder(s). Snowflake may not be associated with, or be sponsored or endorsed by, any such holder(s).