

The SymphonyAI guide to generative AI copilots

How businesses can get the most
out of generative AI

Introduction

Generative AI copilots are everywhere. In the space of a year, AI has become commonplace in the workplace. According to some estimates, the artificial intelligence market within financial services will be worth [\\$85.7 billion by 2030](#). The market is difficult to predict however; this year, financial institutions are expected to spend [\\$5.6 billion](#), which is 14x less. The only points that experts can agree on are: that AI is here to stay, that it's the biggest shakeup of how we work since the Internet, and that implementing AI systems into your business operations will most likely occur sooner rather than later.



AI is the biggest shakeup of how we work since the Internet.

In this guide, we'll be talking about generative AI copilots, which is the term that is already becoming synonymous with the use of gen AI in the workplace.

To help you assess the opportunity copilots could present your organization, we'll be taking you through:

- The questions you should ask
- The research you should undertake
- The implications of using gen AI
- What to expect

If an organization within financial services uses a generative AI copilot effectively, it can increase productivity by up to 70%, enjoying huge efficiency gains thanks to AI generated summaries, recommended actions, and consistent report structures. The aim of this guide is to help your organization get there.

Increase
productivity by
up to 70%
with a generative AI copilot

What is a generative AI copilot?

A generative AI copilot is a generative AI-powered assistant. It helps the user with various tasks by searching text, collating and summarizing information, and writing reports.

Alongside working with text and numbers, generative AI copilots can generate images based on a text prompt, generate code in a variety of programming languages, and search the Internet for relevant articles (that could even be summarized and included in a report).

Using a generative AI copilot is easy for many reasons:

- **Its conversational interface** – just as with chatbots, users type their question into a text field, and the copilot will reply with the information using natural language that can easily be interrogated and queried.
- **It searches and generates text** – Generative AI copilots for business can interact with your existing software, making it simple to use. For example, if working with a lot of numbers, a copilot can quickly extrapolate relevant information and do the math required in the time it used to take for someone to open a spreadsheet. Not only can it do the math, but it could also summarize its efforts and explain its working if necessary.

- **Many plugins are available** – Because of the way publicly available generative AI copilots work, there are limits to business use. Most organizations require a copilot that works specifically for their use case, are trained using the relevant large language models (LLMs), and can “plug in” to existing solutions.

With all these abilities available, it's easy to see how **generative AI copilots are moving from 'optional' to 'essential' in the business world**. They are effective, drive efficiency, and lead to enhanced productivity across teams and departments.

Within financial services, banks are expected to spend [5.6 billion on generative AI in 2024](#). Let's take a closer look at one example of generative AI for this industry, a tool that specializes in helping financial crime investigators – the Sensa Copilot.

Note:

Throughout this document: **'Generative AI copilot'** or **'copilot'** refers to any generative AI copilot **'Sensa Copilot'** refers to SymphonyAI's copilot capability within the Sensa Investigation Hub


What is the Sensa Copilot?

The [Sensa Copilot](#) is part of the [Sensa Investigation Hub](#), a single, enterprise-wide solution for all financial crime investigation needs, powered by generative AI. Detection-agnostic, the Sensa Investigation Hub is configurable to any use case, business unit, or asset class. It offers a unified, entity-centric view of risk and allows for 70% more efficient investigations. Importantly, it is available as an augmentation to allow organizations to start using generative and predictive AI as soon as possible, without immediately having to replace their current technology.

The Sensa Copilot is a key component of Sensa Investigation Hub. It harnesses the power of generative AI to source, summarize, and analyze data, freeing up financial crime investigators to focus on making informed decisions. It can interrogate all relevant sources including internal systems, third-party data, and curated web searches to collect information available on the activity of a suspicious entity.

Generative AI summarizes this information into an easy-to-understand, natural language narrative describing why an alert has been raised and what type of risk your organization is exposed to. Investigators interact with the Sensa Copilot via an intuitive, chat-style window and can use generated narratives to populate case logs and reports in a format that is easily explained to regulators - investigations follow consistent processes and reports are composed in unambiguous, exact language, with every step fully auditable.

In this way, the Sensa Copilot is not only efficient but also effective. It will accelerate investigations, but just as importantly, it will improve overall investigation quality, including disclosures. By having the Sensa Copilot as an AI assistant, organizations can ensure that the investigator is taking recommended actions, seeing the most relevant information using generated narratives, reviewing the most credible web searches, and drafting consistent reports using generated report narratives.



Examples of the Sensa Copilot

Sensa Investigation Hub was developed for financial institutions to help in their combating of financial crime, and the Sensa Copilot is a key component.

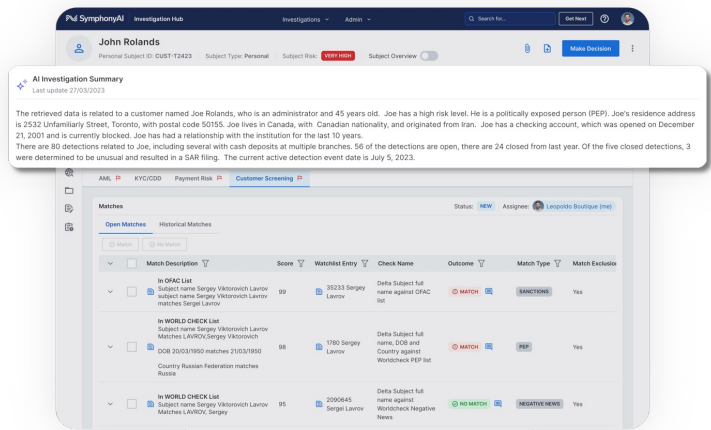
The Sensa Copilot ingests all relevant information on the subject being investigated from every available data source, including third-party detection engines and data sources that are connected to Sensa Investigation Hub. It uses this information to provide an entity-centric view of risk including summaries and key findings to the investigator.

Context-aware, it understands what type of investigation is being undertaken and where the investigator is in the process, calibrating its response to best suit their needs based on previous suspicious activity reports (SARs).

Impressively, the Sensa Copilot can handle almost any risk. From KYC/CDD issues and sanctions screening to AML and payment fraud events, it helps investigators tackle the full range of threats an organization faces, as can be seen in the following examples.

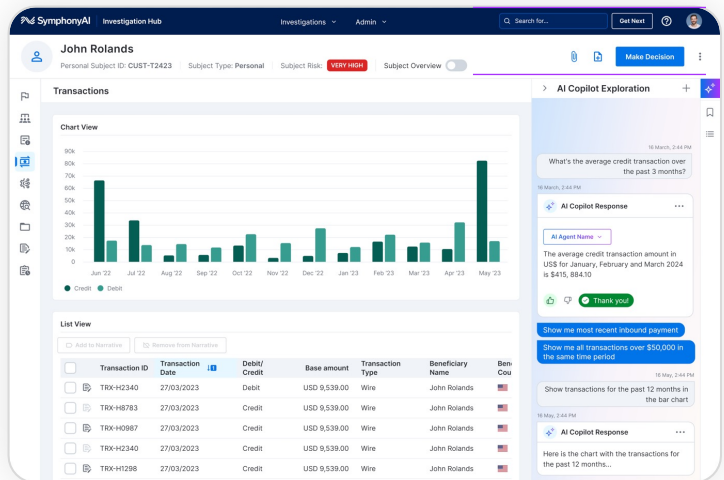
Example #1 – Summarizing a current financial crime investigation

In this example, the Sensa Copilot offers a summary of the case and the ongoing investigation so far. This can be seen in an AI generated paragraph at the top of the screen and will update as the case continues.



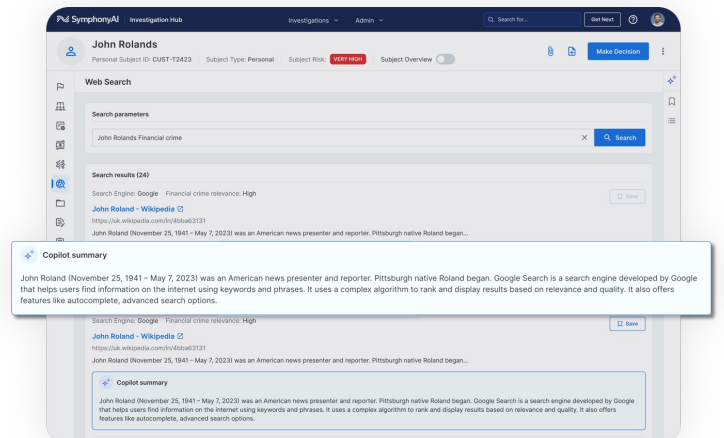
Example #2 – Aggregating and displaying information for an investigator

In this example, there is an ongoing investigation into a customer. The investigator has asked for all the transactions the customer has made in the last twelve months. The Sensa Copilot works with Sensa Investigation Hub to show this information as clearly as possible.



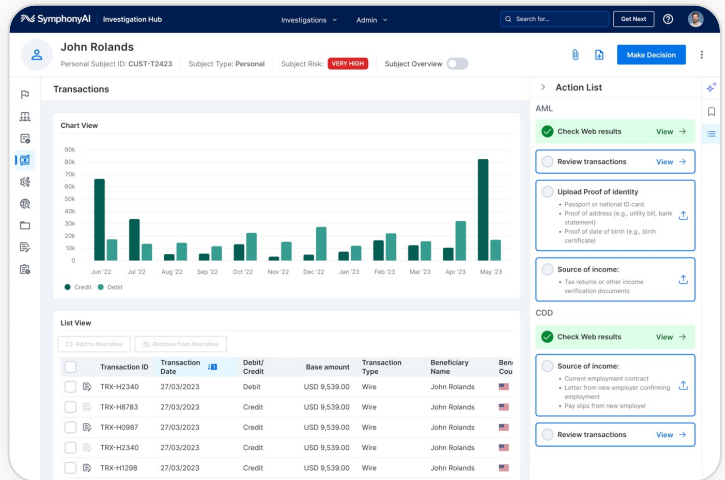
Example #3 – Summarizing web articles for use in an investigation

In this example, the investigator has asked the Sensa Copilot to do a web search on the name of the person being investigated. If the copilot finds a match, these links will be made available to the investigator. For added convenience, the Sensa Copilot summarizes the information, notes its relevance, and can save the article for use in the case narrative. This in turn can be used in a SAR but doesn't have to be.



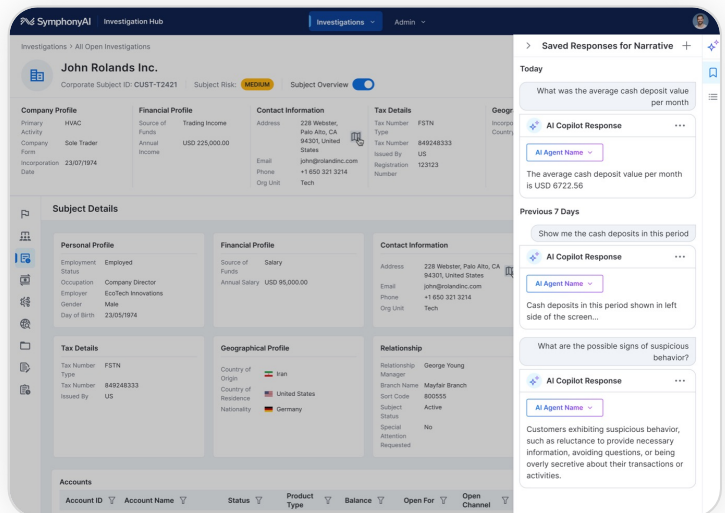
Example #4 – Provide an action list

In this example, the Sensa Copilot helps an investigator with suggestions for improving a case. From offering to review transactions and source of income to checking web results, it's a comprehensive approach to financial crime investigations.



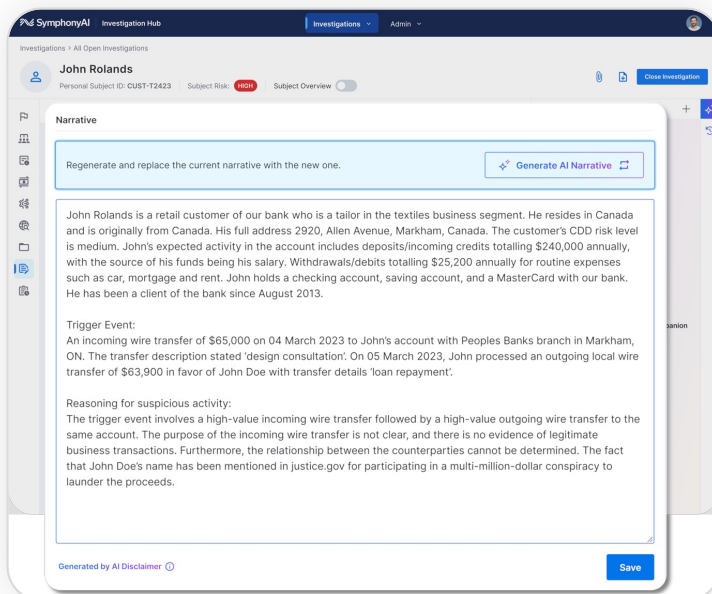
Example 5 – Bookmark key information

In this example, an investigator is reviewing information via the Sensa Copilot, assessing responses bookmarked throughout the investigation so far. Any bookmarks will be added to the narrative summary seen in Example #1 and accessible to any investigators as the investigation progresses – preventing doubling up on efforts.



Example #6 – Creating a SAR from the information collected and collated

Here, all the information that an investigator has collected needs to be written up into a narrative. As new information comes to light, an investigator can refresh the report narrative to include it. The Sensa Copilot generates the full draft text within just a few seconds.



This is just a small selection of what the Sensa Copilot can do within anti-financial crime investigations. Even so, it's clear to see the benefits on productivity, efficiency, and effectiveness for financial institutions.

Key benefits of Sensa Copilot



Subject-centric view of risk

Centralize and connect data from across your tech ecosystem for a truly holistic view.



Gen AI powered investigation

Sensa Copilot enables investigations to be up to 70% faster and more consistent.



Detection engine agnostic

Seamlessly integrate with any third-party detection system or data source.



Achieve total transparency

Completely auditable workflows and superior explainability ensure enhanced compliance.

How to get the most out of a generative AI copilot as a CRO

In terms of SymphonyAI products, it's important to note that financial crime investigations are still very much investigator driven. Investigators are wholly responsible for decisions, but the Sensa Copilot is there to make the process easier and more effective. Though this is the case, there are still topics to be aware of when implementing generative AI copilots within your organization.

It's best to run proofs of concept with potential vendors (though there aren't many currently) to better understand potential data privacy and security concerns, the phases of a copilot's implementation, and how to attain internal stakeholder buy-in. The safety and reliability of using a copilot as well as its impact on the culture of your company are also important to investigate.


Let's address these topics in turn.

Running a proof of concept

Introducing new technologies into your organization is a great opportunity to improve productivity, revenue, and profitability. However, it can be difficult to know the effects of the technology in advance, which is where running a proof of concept (PoC) comes in. Not only is this good for your company in assessing the technology and its capabilities effectively, but it also allows both the customer and the vendor to work together on getting the best fit for an organization's specific requirements. In this way, it is possible to spot any challenges early. Perhaps the generative AI copilot doesn't solve problems as expected or the integration or

infrastructure requires support, PoCs are the best gauge of what to expect from customer support and technological performance, through to expertise, and return on investment.

A final reason to go for a PoC is that it is cost-effective. Running a PoC is brief, and it doesn't require full implementation (and the associated costs) immediately, with both sides benefiting from the situation. A PoC allows for an informed choice in an efficient, timely manner with proven data to back up the decision. It's an important step in ensuring that the solution aligns with your business goals and requirements.



Data privacy and security concerns

The security and privacy implications of using a generative AI copilot within a work environment, especially where it is handling confidential information, is a critical concern for all organizations so it's essential to understand exactly how the copilot works before implementation.

This starts with understanding how a generative AI copilot has been trained, what language learning models (LLMs) may have been used, and to make sure that the information it has access to is in a fenced off environment that can't be accessed by someone that doesn't have access to internal private networks and processes.

With Sensa Copilot, SymphonyAI ensures that data entered by users doesn't leave the client's own environment. As such, there is no

risk of proprietary information leaking into the model and becoming available to other users. The inputs, outputs, models, and training data are only available to authorized internal employees and investigators only have access to customer data for the duration of the investigation session. Customer data is only be used in line with existing data processing policies, and third-party data will only be collected from reputable counterparties that comply with legislation like GDPR and CCPA. Once the session is closed, the data is wiped from the Sensa Copilot.

In terms of training, Sensa Investigation Hub uses Azure OpenAI models, trained on the very best ('gold star'), human generated cases. Customer data can also be used to train models but only for that specific customer.





Phased implementation

Introducing a generative AI copilot into your work environment can be overwhelming, which is why it's best to stagger its release over a set period. This allows users to slowly incorporate its use into their working day, while also allowing team leaders to assess the copilot's effect on the company without going 'all in' on day one.

A common approach to the incorporation of new technology is by thinking in terms of phases. For the Sensa Investigation Hub and the Sensa Copilot, it might go something like this:

1 Planning phase

Presenting Sensa Investigation Hub to those who will be using it, including highlighting the Sensa Copilot. This is the perfect opportunity for beginning to think about how AI will be used within the wider AI strategy of the business, for understanding the specific needs of the team and organization, and for setting goals.

2 System installation phase

Tuning the Large Language Models (LLMs) for narratives and modify database views where needed for required data. This is the information that the Sensa Copilot will use when generating its responses to the investigator.

3 Onboarding phase

Provide training materials and sessions to users, alongside their account details and platform access. It is wise to roll out slowly, perhaps team-by-team, and observe the user experience. A good place to start is with experienced investigators that can road test the new software including the Sensa Copilot.

Based on their feedback, managers can use built in UI (user interface) tools to create investigation workflows that match their needs and policies, taking full advantage of Sensa Investigation Hub's many capabilities.

4 Operation and testing phase

As more users gain access and start to use the Sensa Investigation Hub and the Sensa Copilot, continue observations, highlighting how it is being used throughout the working day.

Ensure that feedback is regularly received and continue fine tuning the approach to maximize the impact, efficiency, and effectiveness of the Sensa Copilot.

SymphonyAI deployment teams go further than these phases, viewing any deployment of a product as a collaborative endeavor. Dedicated to helping organizations succeed, we offer personalization, gap and workflow analysis, tech assessments, and training to ensure that financial institutions have key concerns addressed as soon as possible, and that they are using solutions – Sensa Investigation Hub and all its capabilities including the Sensa Copilot – to their full potential.

Stakeholder buy-in

When implementing a generative AI copilot into your business, it is crucial to receive the support, commitment and alignment of key stakeholders. This will ensure a smooth execution of the plan and the ultimate success of the project.

To achieve stakeholder buy-in with a copilot, it's first necessary to align it to the overall AI strategy of an organization, understanding and connecting it to the long-term goals of the company, and the roles financial intelligence units (FIUs) have in the overall vision.



To achieve stakeholder buy-in with a copilot, it's first necessary to align it to the overall AI strategy of an organization.

Alongside this, identify all relevant colleagues that will be impacted by using it. Understand their perspectives by addressing both their interests and concerns and take the time to communicate clearly and frequently as the project evolves. One approach would be to keep everyone in the know as each phase of implementation occurs.

To get people excited about using the generative AI copilot, it's vital to explain the many benefits it will deliver and how the stakeholders can contribute to company goals, providing as much data as possible to support the position. If there is pushback, be sure to listen and collaborate with the stakeholders where possible; involve them in planning sessions, incorporate their ideas, and seek their input where required.

Finally, allow the generative AI copilot and its results to speak for itself. Delivering on the promised outcomes will build trust among your team members and make the organization stronger, creating a supportive environment ahead of the next big technology rollout in the future.



Safety and reliability

As with all technology, there is a safety and reliability aspect to incorporating AI into a business. It is vital to ensure that a chosen generative AI copilot can handle the data it is given and offer answers that are unbiased and devoid of information that may have come from unintentional leading questions. Using a copilot to summarize meetings for example, may see misinterpretations or a complete omission of information that those present may feel important to the situation being discussed.

Sensa Copilot mitigates the issue of omissions in AI-generated narratives by having the investigators select the information that is included from their research; which transactions, detections, web searches, etc., to include. This is known as a context injection. Due to the volume of research that could be added to a summary, the Sensa Copilot may provide its own summarization (e.g., 'several websites with suspicious information were found.')

Another common concern is hallucinations, meaning where a generative AI copilot makes up facts. Although an issue generally,

it is possible to manage and control this in a business use case. For example, one way Sensa Copilot addresses the issue is by using retrieval augmented generation. This ensures that the narrative is based on information retrieved from databases rather than on probabilities.

When first implementing a generative AI copilot, it is always best to have a human reading through AI output to ensure that nothing important is missed. Generative AI can increase efficiency but it's ultimate success will be down to how human's get the most out of it.



When first implementing a generative AI copilot, it is always best to have a human reading through AI output.



Cultural impact

Workplace dynamics and culture are a very important part of office life. Introducing a generative AI copilot into proceedings can shift those team dynamics and company culture unless steps are taken to mitigate any possible problems before they occur.

For instance, generative AI is often cited as achieving time savings for users. Financial institutions with the Sensa Copilot have noticed productivity improve by as much as 70%. That's a significant saving for any organization as it allows for focus on more important work. In this vein, organizations should set clear benchmarks and targets to assess benefits, while training and feedback loops are essential to develop a generative AI copilot that suits the needs of an organization.

With the rise of gen AI, there has also been some understandable pushback regarding the ethical implications of such technology. With a generative AI copilot's ability


to gather data, write reports, and suggest detailed plans, there are concerns over the changing nature of work. It is important to address these ethical issues and make sure that your organization is on top of them.

One such area might be exclusively depending on an AI copilot, which diminishes or erodes the skills of employees who now might not have to think about typing emails or drafting documents. To ensure organizations are getting the most from both AI and employees, keeping humans in the loop is essential.

Used effectively, a copilot can improve productivity and efficiency. However, it is not a replacement for collaboration, creativity, and the improvement of interpersonal skills, which all go towards creating the inclusive workplace dynamic and culture that will have been initiated before the advent of a generative AI copilot.



A copilot is not a replacement for collaboration, creativity, and the improvement of interpersonal skills.



Frequently asked questions

What is the best use of a generative AI copilot?

The best use of a generative AI copilot is the use that will see your company becoming more effective, productive, and efficient. There are so many different use cases that it is impossible to list them all.

The Sensa Copilot has been designed to be most useful for financial institutions fighting anti-financial crime. As such, it can be used for teams working in AML transaction monitoring, sanctions screening (including PEP and adverse media), payment fraud, and KYC/CDD.

How do you interact with a generative AI copilot?

As seen in the examples earlier in this document, a generative AI copilot acts similarly to a chatbot, using natural language dialogue. A user simply types in their query (or 'prompt') and receives an answer using the natural language of their choice. Unlike using a free product, a generative AI copilot for business use is drawing on a specialist set of training data – e.g., financial crime investigation to guide users within on the context of the specific use case.

A copilot can also be used to automatically generate summaries or narratives. In this instance, such as with the Sensa Copilot, it might only require a click of a button to generate the report that you need. Naturally, this depends on the capabilities of the generative AI copilot that you are using.

Can you delete part of a question or prompt once it has been asked?

This is not currently possible with most generative AI copilots for business use. To go in a different direction with your questioning, it is best to start a new interaction with it.

How do you install and implement a generative AI copilot?

This depends on the copilot. The Sensa Copilot is only available as part of Sensa Investigation Hub and comes as part of the package, which is available on the [Azure Marketplace](#) or [directly from SymphonyAI](#).

The Sensa Copilot has been designed to be most useful for financial institutions fighting anti-financial crime.

How long does it take to implement a generative AI copilot?

Timelines vary depending on multiple factors, but it is best to think in terms of months rather than weeks to ensure that the generative AI copilot is best placed to work with the data that you are providing it with and that everybody within your company is suitably aligned. It is also best not to rush these things due to a multitude of factors mentioned in previous pages (cultural impact, security, etc.).

What data is a generative AI copilot trained on?

Generative AI copilots for business use vary significantly from free models such as OpenAI's ChatGPT or Meta's LLaMa.

Business copilots such as the Sensa Copilot use a variety of LLMs that have been trained on proprietary data to ensure that they are the best generative AI copilot available for their specific industry use. The Sensa Copilot ingests data from internal and third-party sources to empower investigators with context information and enhance understanding of entities and alerts. It is not advisable to use a copilot that has not been trained in this way for a variety of reasons including data security and privacy concerns, alongside a lack of control over AI output.


While interacting with the data that an organization provides, the Sensa Copilot can learn how better to administer and use your data most effectively. It is, in a sense, always learning and improving.

How can an organization take advantage of generative AI like a copilot without risking security and privacy?

Every generative AI copilot on the market will handle this area differently. For SymphonyAI, security and privacy are paramount, and all products are built to uphold the highest standards.

The generative AI functionality of the Sensa Copilot is deployed inside an organization's existing environment and controls. The inputs, outputs, models, and training data are only available to authorized internal employees. Customer data will only be used in line with existing data processing policies, and third-party data will only be collected from reputable counterparties that comply with legislation like GDPR and CCPA.

Critically, staff outside of the customer (e.g., vendor teams) will not be able to see any of the customer data as it is transferred via API. Investigators will have access via the Sensa Copilot to customer data for the duration of the investigation session. Once the session is closed, the data is wiped from the Sensa Copilot. The only data the Sensa Copilot stores is strictly for training (e.g., "gold star" use cases) and required retrieval augmented generation.



About SymphonyAI

SymphonyAI, 2024 Microsoft Partner of the Year for Business Transformation – AI Innovation, is building the leading enterprise AI SaaS company for digital transformation across the most critical and resilient growth verticals, including retail, consumer packaged goods, finance, manufacturing, media, and IT/enterprise service management. SymphonyAI verticals have many leading enterprises as clients. Since its founding in 2017, SymphonyAI has grown rapidly to 3,000 talented leaders, data scientists, and other professionals. SymphonyAI is an SAIGroup company, backed by a \$1 billion commitment from Dr. Romesh Wadhvani, a successful entrepreneur and philanthropist.

SymphonyAI provides AI-focused SaaS solutions developed from the ground up for fighting financial crime. An end-to-end offering compiled from more than 25 years of knowledge in tackling money laundering and fraud, our award-winning product ecosystem comprises an innovative and modern approach to financial crime investigation using world-leading predictive and gen AI. From KYC/CDD and transaction monitoring through to payments fraud, entity resolution, and sanctions, PEP, and adverse media screening, SymphonyAI is a trusted name in finance technology with more than a third of the world's top 100 banks enjoying the power, efficiency, effectiveness and productivity that our solutions provide.

Learn more at the [SymphonyAI website](#).