SymphonyAI

From Theory to Action:

AI Agents Transforming Financial Crime Compliance

Frequently Asked Questions

During the webinar, attendees submitted questions covering data security, AI training, system integration, deployment architecture, and product roadmap. The following answers address those questions directly, providing the detail needed to evaluate agentic AI in the context of your institution's financial crime and compliance programme.

Want to continue the conversation?
Contact your SymphonyAI representative or send an email to fs.connect@symphonyai.com to learn more about our agentic AI capabilities.

Data and training

Q1. What kind of data is used to train the AI agents, and where does it come from?

Training depends on the agent type. Initial models are built using synthetic data curated by SymphonyAI, based on over 25 years of sanctions expertise, and are purpose-designed to protect customer information during development.

Once deployed in your environment, agents train continuously on your own data, responses, and interactions. The result: agents that remain globally consistent in design and capability but become increasingly specific to your institution over time.

Q2. How can institutions be confident that the AI is working with clean and reliable data?

Data quality is managed at the point of ingestion. The platform includes validation processes to flag incomplete, inconsistent, or anomalous inputs before they reach the agents.

Institutions retain full visibility into the data sources being used, so quality concerns can be identified and addressed at the source and not discovered downstream in an investigation.

Customisation and market adaptability

Q3. Can the STR/SAR module be customised to reflect different regulatory requirements?

Yes. Regulatory reporting formats and content requirements vary by jurisdiction, and the platform is built to accommodate that. STR modules are configured market by market.

As an example, a report built for Malaysia will have a different format and content structure from one built for Australia and both will align to the respective regulator's expectations.

Consistency and audit trail

Q4. How does the platform ensure reproducibility of narratives and insights over time?

The platform maintains consistency through model versioning and a structured audit trail for every AI-generated output. Decisions, narratives, and case findings are logged and fully traceable.

Infrastructure and deployment

Q5. Is the platform hosted on Google Cloud, AWS, or another provider?

The platform is cloud-agnostic. Deployment infrastructure is determined by your institution's preference whether that is a specific hyper scaler, a private cloud environment, or an on-premise configuration.

Microsoft is our global key partner, however SymphonyAI works with your existing infrastructure strategy, not against it.

Q6. Banks operate under strict data privacy obligations. What measures are in place to prevent data leakage when using AI software?

Data security is built into the deployment architecture. The platform operates in a closed environment, owned and controlled by the bank. No data is transmitted to or shared with external parties.

Q7. Do you offer an on-premise deployment option for institutions subject to data domicile requirements?

Yes. For institutions operating in jurisdictions where central banks or regulators require that data remain in-country, SymphonyAI supports on-premise deployment. The platform architecture is designed to accommodate strict data residency requirements without compromising on functionality.

Data integration

Q8. Can the AI access data from multiple systems, or is it limited to a single data source?

Multi-source data integration is supported. How data is ingested and which systems the agents interact with is determined during implementation, in alignment with your institution's data governance policies.

The platform is designed to comply with each bank's data access controls meaning the scope of integration is defined by your IT and compliance teams, not imposed by the vendor.

Investigation workflow

Q10. The investigation hub appears to add steps to the process. Will this increase complexity for investigators?

The intent is the opposite. The platform fundamentally reframes the investigation process moving away from the traditional L1/L2/L3 structure toward a more efficient, intelligence-led model.

Two factors drive this: the process itself is more streamlined, and complementary detection capabilities significantly reduce low-relevance alerts (false positives). What appears as additional steps on screen translates to less manual effort, deeper risk analysis, and faster case resolution overall.

Product scope and roadmap

Q11. Can the platform support capital markets intermediary products such as stockbroking, futures trading, and non-account-based advisory services?

Capital markets coverage is on the product roadmap. The current platform is focused on banking and financial crime compliance use cases, with planned expansion into capital market intermediary products in future releases.

If this is a priority for your institution, we welcome a conversation to understand your requirements and share our roadmap in more detail.

Q12. Can institutions choose which AI or LLM they want to leverage, or are they limited to Eureka AI?

Eureka is SymphonyAI's proprietary vertical AI stack, purpose-built for enterprise financial services use cases and optimised for the complexity of financial crime compliance.

The platform supports integration with other large language models for Co-Pilot and similar functions. Customers are not locked into a single LLM, and SymphonyAI works with institutions to align AI deployment with existing technology preferences where appropriate.

Integration and onboarding

Q13. Does the Co-Pilot require system integration? Does it need to be trained before it can generate case recommendations?

The Co-Pilot is pre-trained and ready to deploy. System integration is required to connect it to your institution's data environment, but no manual training is needed before it can generate recommendations.

From deployment, it learns continuously from your institution's usage, interactions, and case outcomes becoming more tailored and effective over time.