SymphonyAI

Agentic automation Q&A

Everything you need to know about Agentic AI automation

Agentic AI guide

Introduction

Agentic AI is the latest form of automation that has everybody excited within the technology world. But what is agentic automation and how can it be used in financial crime prevention?

Eric Murray, Director – Product, Generative AI within SymphonyAI's financial services division, explains everything you need to know.

Eric is responsible for integrating and expanding the use of GenAI across the product suite. He is an engineer with extensive experience in software development, data science and analytics, and customer-facing delivery roles. Previous roles include pre/post sales, technical director, and customer success, solution engineering, and developing custom solutions for numerous defense and space agencies.

What’s APA?

What is agentic process automation (APA) and why do you see it as the future of AML compliance?

Eric Murray: Agentic Process Automation or Agentic Workflows are system components that operate on unstructured data in scenarios traditionally reserved for humans. These scenarios are simply too complex for regular software systems. Agentic flows use Large Language Models (LLMs) as reasoning engines and dynamically decide which operations to execute and which systems to interact with based on natural language instructions rather than complex algorithms.

Currently, a large part of an investigator’s time is spent gathering and collating relevant and required data from different systems based on nuanced understanding of financial crime alerts, customer details from numerous internal and external systems, and their organizations risk and compliance rules. Agentic automation can instead perform much of this research, data aggregation, and relevancy decisions for these complex tasks. Agents can deliver an investigation package and recommendations to the investigator for a quick and simplified decision-making process.

APA vs. RPA

What is the difference between agentic process automation and robotic process automation (RPA)?

Eric Murray: RPA is ideal for automating repetitive tasks with clearly defined rules and structured inputs. RPA would be useful for extracting fields from a standardized form and entering them into a database. However, it struggles with unstructured data or tasks that require adaptation or decision-making.

Agentic process automation goes further. Agents can handle unstructured inputs, reason through ambiguity, and adapt to variations. For instance, an agent could read a free-form email, extract the relevant information, and populate a database, even if the format or phrasing changes while also improving over time through learning.

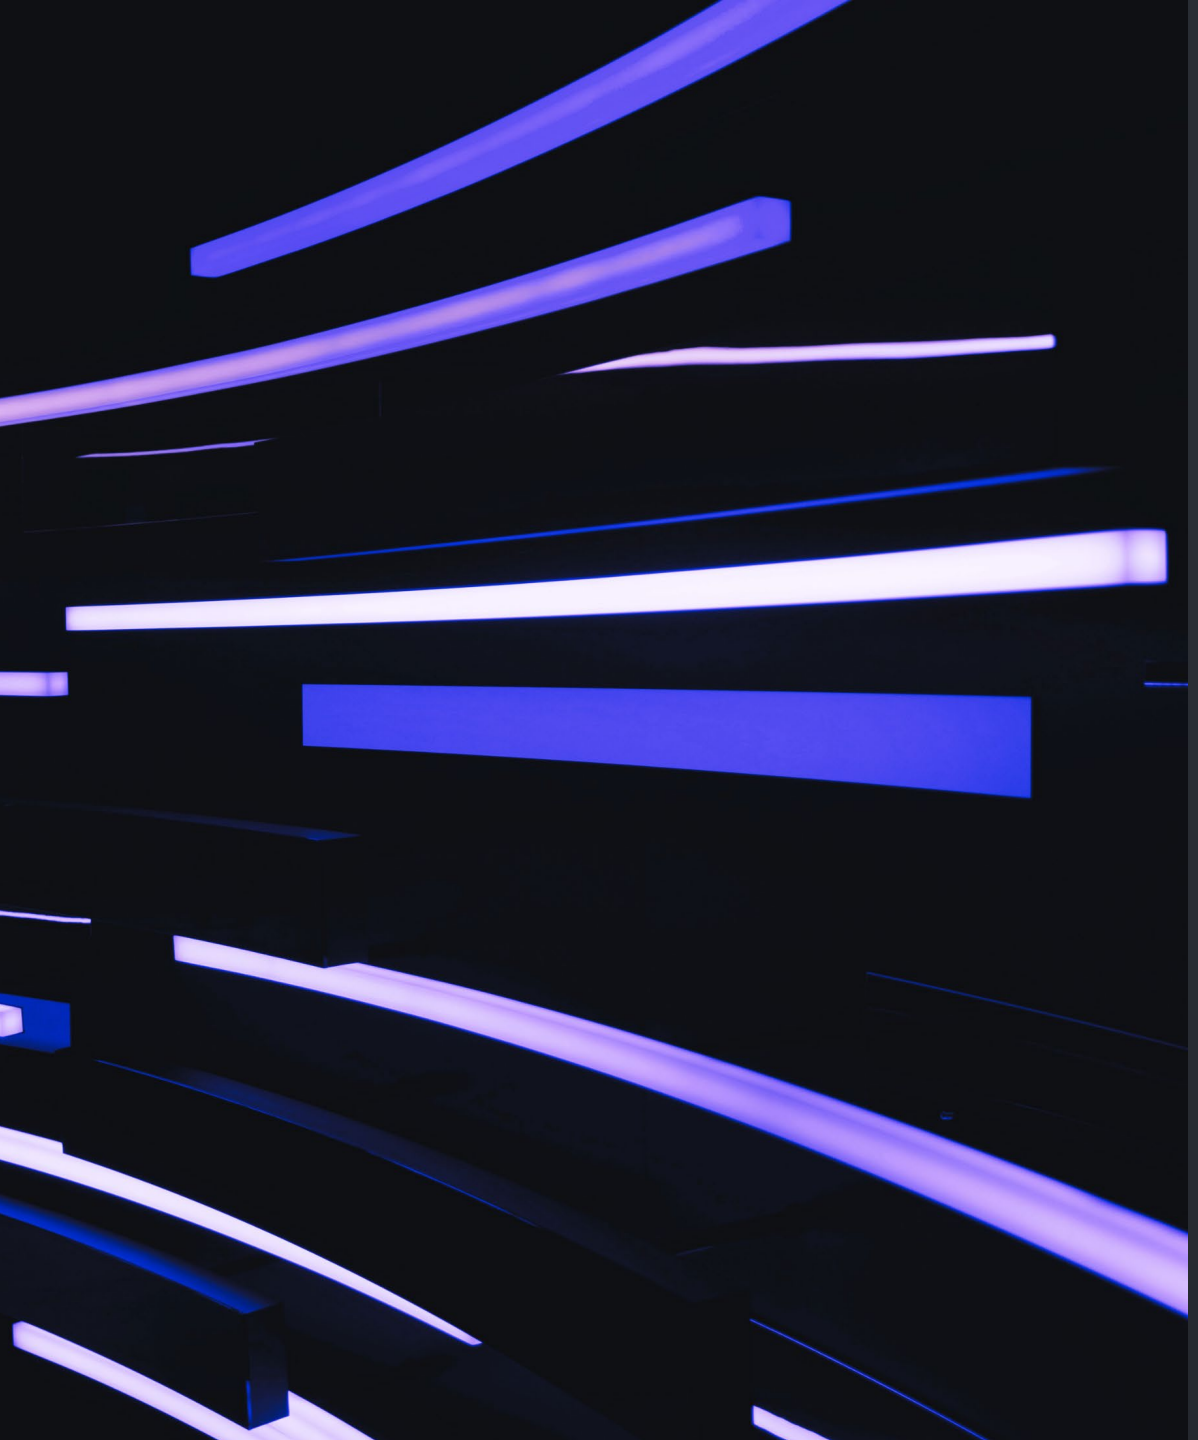

Responsible AI

What guardrails need to be put in place when developing autonomous agents?

Eric Murray: Autonomous agents require strong guardrails; full autonomy isn't viable in most enterprise environments. While it's harmless for an agent to research and draft a blog about trees, it's a different matter when agents process sensitive data, write code, or engage with customers.

At SymphonyAI, we follow a company-wide Responsible AI framework that governs both predictive and generative models. This framework emphasizes accountability, transparency, reliability, safety, privacy, and security.

For agentic systems specifically, we've found that limiting full autonomy is key. Rather than letting agents operate unchecked, we use orchestration layers to coordinate smaller, purpose-driven agents. This allows us to:

- Monitor tool use, agent steps, and outputs more easily
- Introduce Human-in-the-Loop (HITL) oversight to build trust and provide feedback
- Boost task success rates and stability

This approach blends control with flexibility, allowing us to deliver robust performance while ensuring agents remain aligned with enterprise expectations.

What’s the difference?

How does agentic automation differ from other AI-driven automation technologies in terms of scalability and adaptability?

Eric Murray: APA differs from traditional AI automation by offering broader adaptability and scalable autonomy.

Most AI-driven automation, like product recommendation engines or image recognition in quality control systems, is designed for narrow, predefined tasks within a specific domain. For example, a recommendation engine predicts what a user might like based on past behaviour, and a computer vision model will detect defects on an assembly line. These systems perform well but operate within fixed parameters and typically require retraining or redevelopment to handle new tasks.

APA is built to manage complex, multi-step workflows across domains, make decisions in real time, and adapt to evolving inputs with minimal human intervention. An agentic system could, for example, coordinate an entire KYC onboarding process from extracting details from unstructured documents, updating systems of record, and requesting clarifications from either bank staff or the customer, learning and improving as it goes.

This flexibility makes APA highly scalable. Instead of rebuilding systems for every new use case, organizations can extend or reconfigure agent behaviours, thus unlocking broader automation with less overhead.

The role of machine learning

What role does machine learning play in enhancing the capabilities of autonomous agents, and how do you handle the data required for this?

Eric Murray: Machine learning (ML) and agentic automation work best together, each serving distinct but complementary roles.

ML models are ideal for structured data and predictive tasks like classifying transactions, scoring risk, or detecting anomalies. These models can serve as the first step, triggering AI agents to take over with human-like reasoning or multi-step workflows traditionally reserved for humans. For example, if a model flags a suspicious transaction, an agent can investigate by gathering additional context, interacting with internal and external systems, and drafting a decision narrative or report with sources, just like a human investigator would.

There's also a feedback loop: the outcomes of agent investigations can be fed back to improve model training and performance over time.

Agents can also help humans interpret model outputs. Users often struggle to understand technical explanations like SHAP values. Instead of showing raw numbers, agents can translate them into clear, actionable insights using natural language the investigators understand. This capability improves both transparency and trust in the agent's decisions.

Agents can also build or use lightweight models on the fly to support their workflows. They can use clustering to detect patterns or isolation forests to identify outliers within structured data. This gives them the flexibility to adapt their analysis as the situation evolves.

Which metrics matter?

How do you measure the success of an autonomous agent once it is deployed?

Eric Murray: Measuring the success of agents or agentic flows requires a combination of outcome-based and operational metrics. At a high level, we look for successful task completion - did the agent finish the job and produce the correct result? To evaluate this, we log all agent activity, including prompts, tool use, completions, and outcomes. From these logs, we track key metrics such as:

- Task Completion Rate
- Processing Time
- Accuracy of Outputs
- Tool Call Success Rate

Our platform also includes fault tolerance for long-running agent workflows. If a tool or step fails, we apply retry logic, i.e., waiting, reattempting, and escalating based on defined rules. This helps improve reliability and raises overall success rates.

When agents operate with Human-in-the-Loop oversight, we also monitor acceptance/rejection rates of agent outputs. Over time, these rates provide insight into agent learning and trustworthiness, ideally showing increasing alignment with human expectations as feedback is incorporated. In short, agent success isn't measured by completion alone. Success is about completing the right task, the right way, and continuously improving through data and feedback.

Find out more about agentic AI automation and how it can improve your processes

SymphonyAI is the leading SaaS provider of AI-powered FinCrime prevention and compliance solutions. We transform our customers' ability to manage risk with predictive, generative, and agentic AI innovation that detects and prevents even the most complex criminal activity.

For more information or to request a consultation, please visit **symphonyai.com**

SymphonyAI