Microsoft

Vetting Transformation: A Strategic Framework built using Agentic AI

October 2025

Contents

Executive Summary

This white paper details an AI-driven transformation of Microsoft's partner vetting process, achieving unprecedented gains in speed, accuracy, and efficiency. Presented by **Akhil Singhal** at the AGMA Conference, the initiative introduced Agentic AI (autonomous intelligent agents) to replace slow, manual checks with dynamic automation. Key outcomes include onboarding time reduced from days to minutes, over 150,000 hours saved in a year, and fraud detection accuracy also improved significantly. The paper outlines the legacy challenges, the continuous improvement methodology (Plan-Do-Check-Adjust cycle), the Agentic AI framework and technical architecture, and the measurable business impact. This project demonstrates original contributions and thought leadership in AI application, setting up a new industry standard for vetting and in similar fraud scenarios.

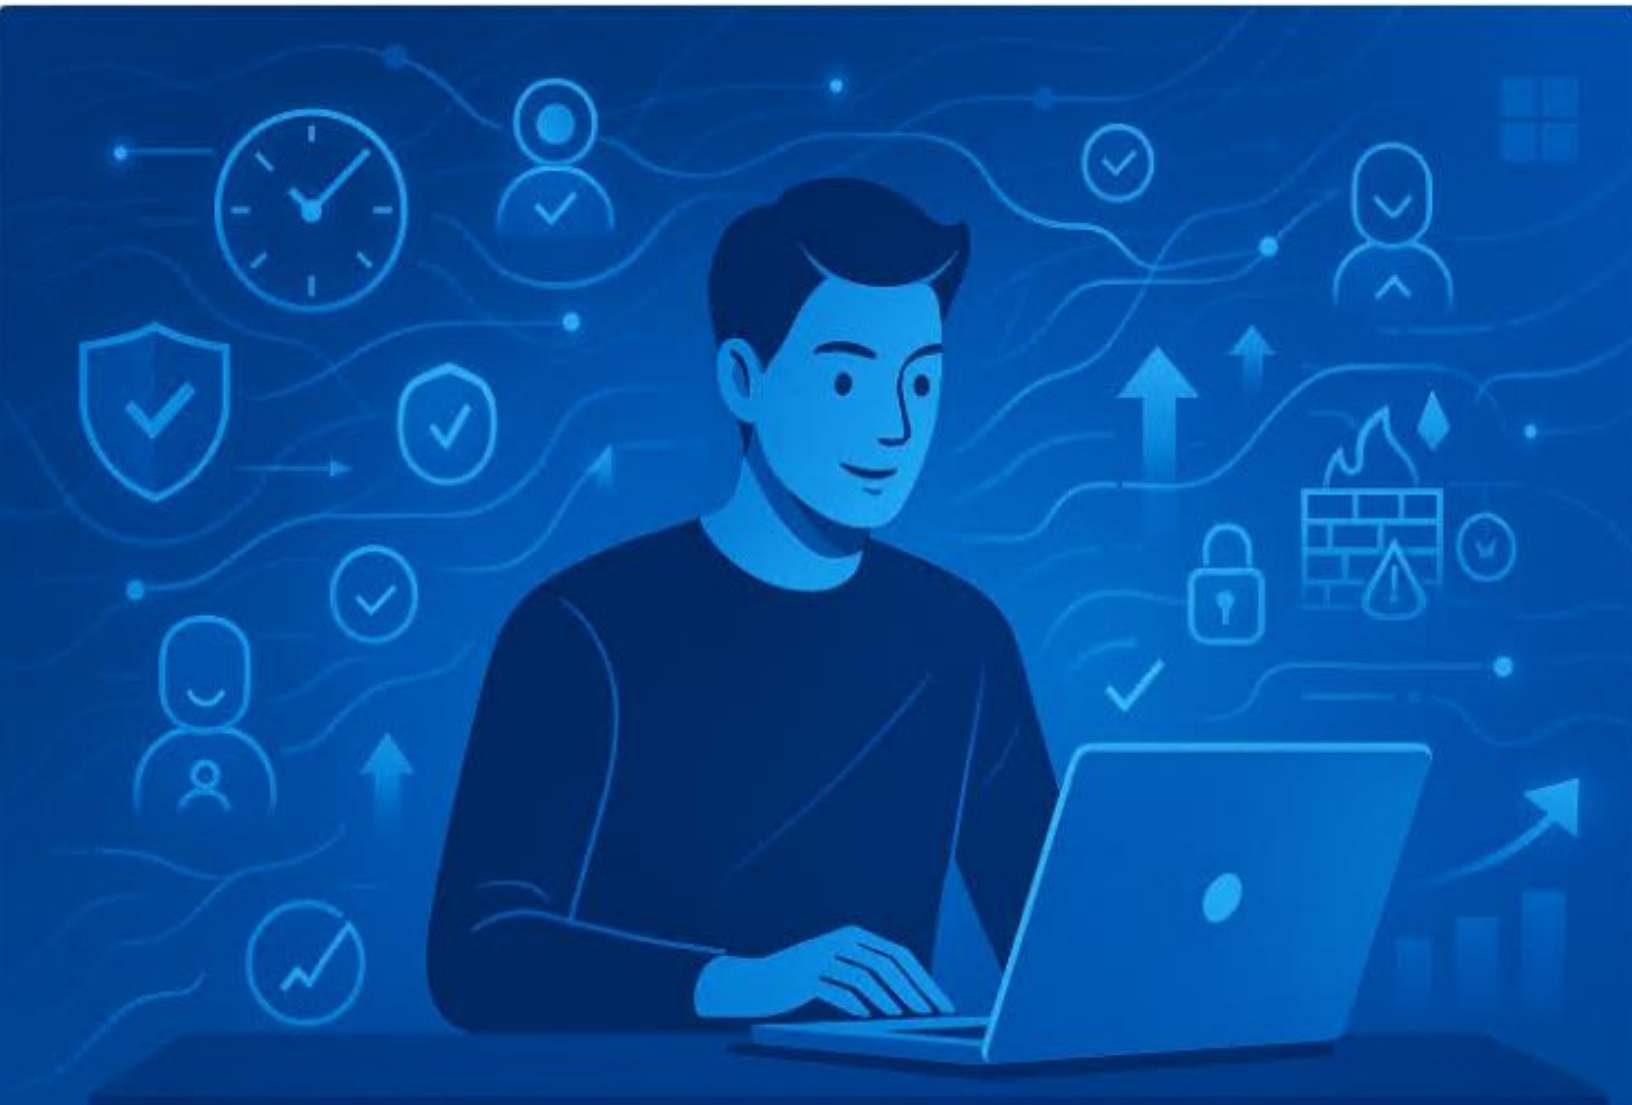

Introduction & Context

In today's digital world, fraud is a pervasive threat that jeopardizes the security, trust, and reputation of individuals and organizations across every industry and region. No sector is immune—fraudulent access to online services affects software, hardware, hosting, telecommunications, and more. Microsoft faces these challenges as well. When unauthorized actors gain access to our products and services, it can harm our customers, partners, and the company itself.

Threats come from a range of adversaries, including nation-state actors, organized crime groups, and individual fraudsters. These attackers exploit vulnerabilities in our systems, programs, promotions, and services. They often leverage stolen resources—such as Azure capacity, licenses, products, credentials, and payment instruments—to carry out cyberattacks, commit financial crimes, violate human rights, or resell assets for profit.

As a leader in technology and innovation, Microsoft is dedicated to preventing its products and services from being misused for fraudulent purposes. Through robust safeguards and ongoing monitoring, we strive to protect our resources from malicious actors and prevent them from falling into the wrong hands.

Entity vetting is a critical process to ensure Microsoft's partners are legitimate, compliant, and trustworthy. Every prospective partner must undergo verification before onboarding and continuous compliance checks thereafter. Traditionally, this vetting involved static rules and manual reviews by the Central Fraud & Risk Abuse (CFAR) team – a group tasked with risk detection, compliance enforcement, and fraud prevention at Microsoft. CFAR's mission is to protect Microsoft and its partners from financial or reputational harm by combining human expertise with AI/ML for thorough evaluations.

Central Fraud and Abuse Risk (CFAR)

Central Fraud and Abuse Risk (CFAR) at Microsoft is a dedicated organization focused on protecting the company, its customers, and partners from fraud. By combining advanced AI models with both real-time and offline risk assessments, CFAR delivers scalable and reliable fraud protection across key customer and partner interactions. The platform integrates data from internal and external sources, enabling rapid detection and a coordinated response to emerging threats. This approach ensures that Microsoft remains a trustworthy business partner, with governance and security controls that adapt to evolving risks and regulatory requirements.

CFAR's mission is supported by a cross-functional team spanning engineering, security, policy, and legal disciplines. This collaborative structure sets a high standard for fraud prevention and operational excellence within the technology industry. Through continuous improvement and innovation, CFAR not only minimizes financial and reputational harm but also helps Microsoft lead the way in safeguarding the digital ecosystem.

Background: The Need for Transformation

In the legacy partner vetting process, the majority of applicant cases—approximately 80%—were handled through automated, rule-based workflows. These static workflows were designed to quickly process straightforward applications that met predefined criteria, allowing for efficient onboarding without human intervention. However, the remaining 20% of cases presented complexities or anomalies that could not be resolved by automation alone. These cases were routed to a manual review queue, where human analysts took over.

Manual review required a meticulous, step-by-step investigation. Reviewers were responsible for verifying the authenticity of submitted documents, conducting thorough compliance checks against regulatory and internal standards, and cross-referencing information across multiple data sources. This often involved checking external databases, contacting applicants for clarification, and collaborating with other teams to resolve outstanding issues.

These pain points underscored the limitations of the legacy system and highlighted the urgent need for a more scalable, accurate, and efficient approach to partner vetting.

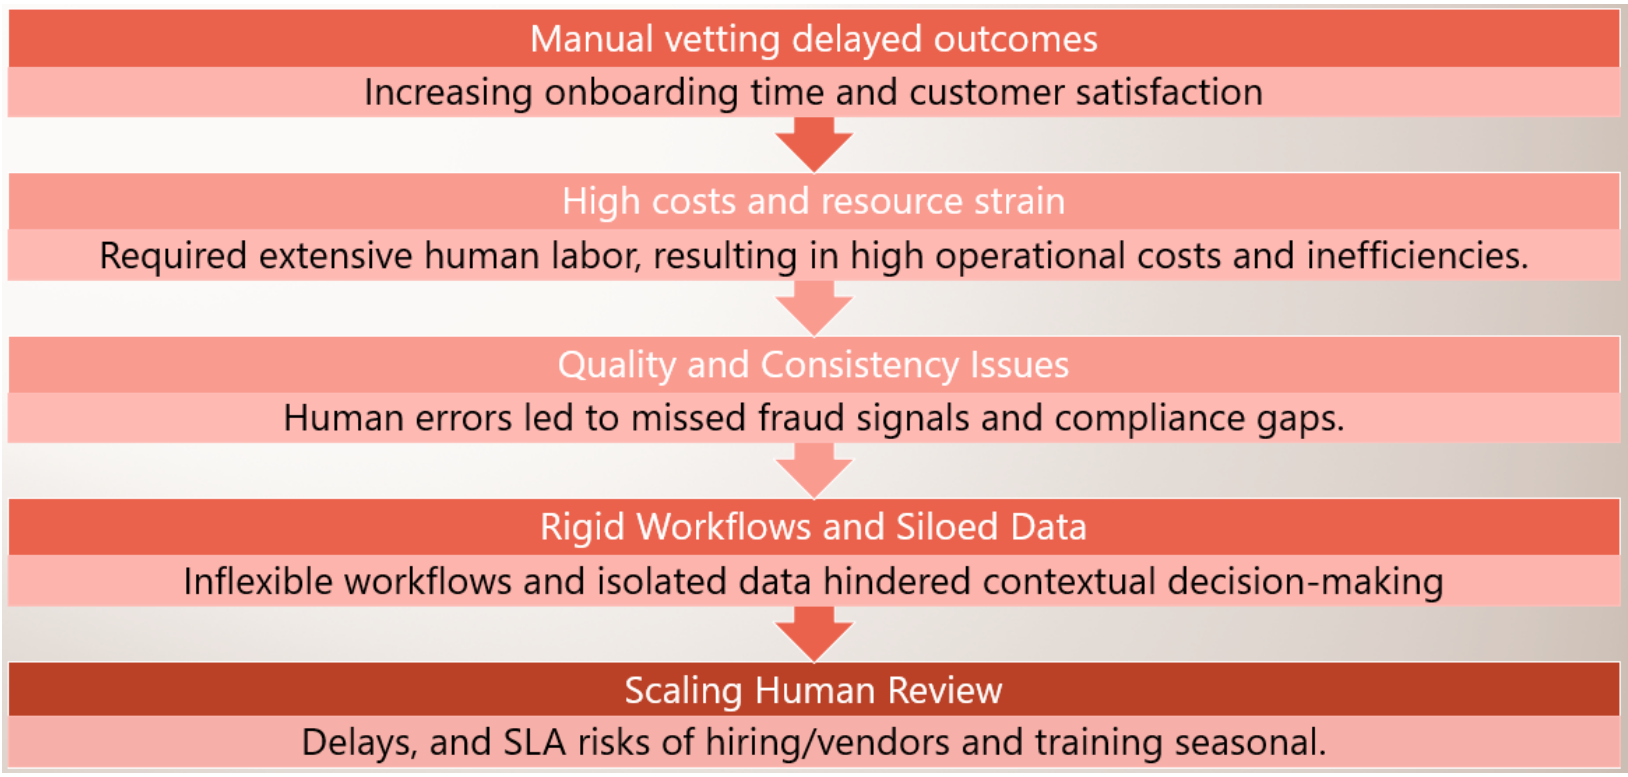

Manual Approach in Identifying Fraud

The manual review process for identifying fraud is inherently complex, involving a series of sequential tasks that demand careful attention and collaboration among reviewers. Each case requires reviewers to meticulously verify documents, perform compliance checks, and cross-reference information from various sources. Some of these tasks are conducted online, leveraging digital tools and databases, while others necessitate offline procedures such as phone calls, physical document inspections, or consultations with subject matter experts. This mix of online and offline activities often results in lengthy feedback loops and delays, as information must be gathered, validated, and communicated across different channels.

Collaboration is a critical aspect of the manual review process, as not all reviewers possess expertise in every area required for comprehensive vetting. To ensure thoroughness, reviewers frequently work together, sharing insights and dividing responsibilities based on their individual strengths. While this collaborative approach enhances the quality of fraud detection, it also introduces additional layers of complexity and extends the time required to reach a final decision. On average, each task within the manual review process takes more than 20 minutes, and in certain scenarios—especially those involving ambiguous or high-risk cases—the overall process can stretch over several days before a resolution is achieved. This extended timeline can impact operational efficiency and delay onboarding or case closure.

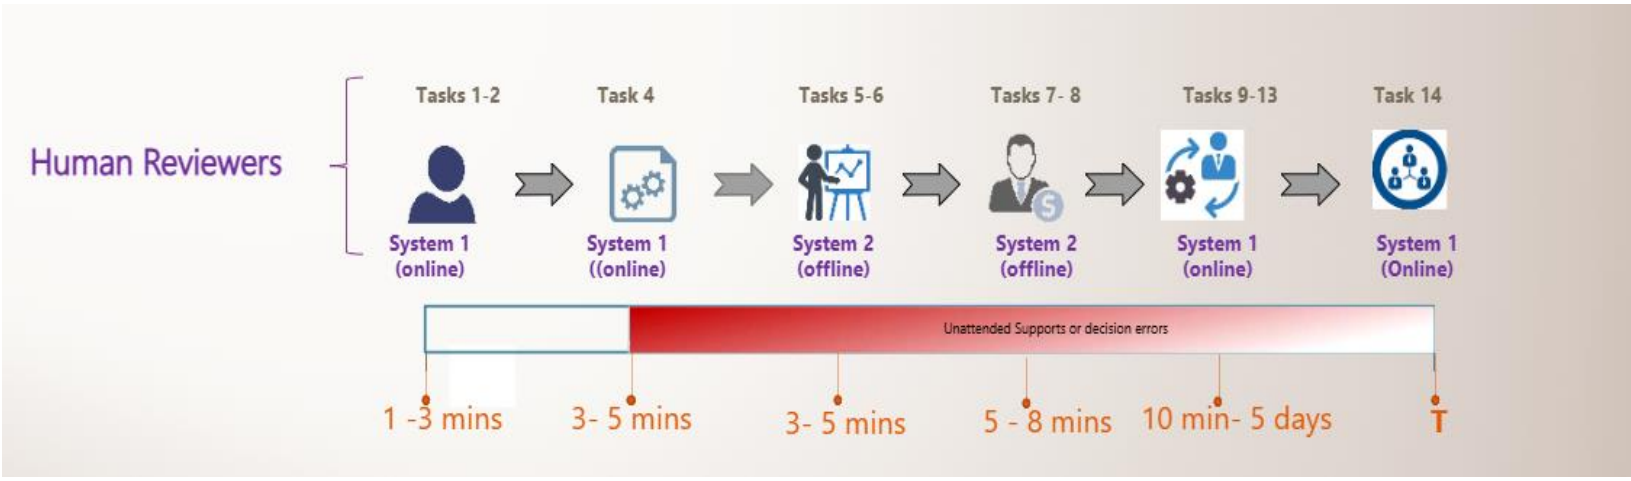

Transformation Methodology: A Continuous Improvement Approach

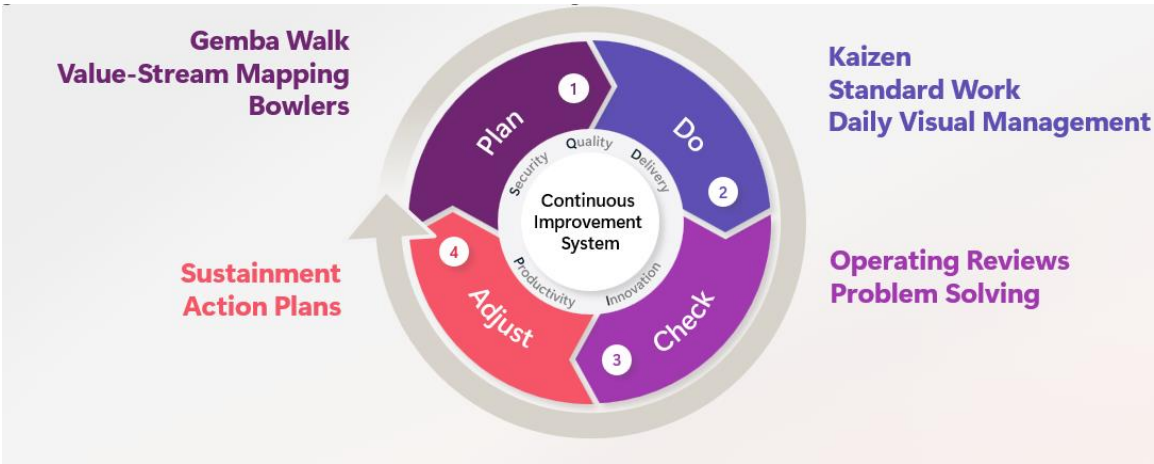

The transformation followed a structured continuous improvement methodology:

- **Plan:** Gemba Walks and Value-Stream Mapping identified inefficiencies and improvement opportunities. Performance indicators were defined in collaboration with human reviewers.
- **Do:** Kaizen philosophy was applied for incremental improvements, establishing best practices and daily performance visibility.
- **Check:** Regular reviews of identified issues and implemented solutions.
- **Adjust:** Ensured improvements were sustained and action plans followed.

This cycle repeated iteratively, driving ongoing gains in accuracy and efficiency.

What is an Agentic AI?

Agentic AI refers to artificial intelligence systems designed to act as autonomous agents—capable of planning, executing, and adapting multi-step tasks independently, often in dynamic and complex environments. Unlike traditional AI, which typically follows static, rule-based algorithms or single-task automation, Agentic AI can:

- **Plan and execute complex workflows:** It breaks down large tasks into smaller, manageable steps and determines the optimal sequence for execution.
- **Adapt to new information:** Agentic AI can adjust its actions based on real-time data, feedback, or changing requirements.
- **Collaborate with humans and other agents:** It can work alongside human reviewers, escalate cases when need, and coordinate with other AI agents for more holistic decision-making.
- **Learn and improve:** Through continuous feedback and telemetry, Agentic AI refines its processes and becomes more effective over time.

Key Characteristics of Agentic AI

- **Autonomy:** Operates with minimal human intervention.
- **Proactivity:** Anticipates needs and initiative.
- **Context-awareness:** Understands the broader context of its tasks and adapts accordingly.
- **Multi-modality:** Can process and integrate information from various sources (text, data, documents, etc.).

Traditional AI vs. Agentic AI – A Comparison

Aspect	Traditional AI	Agentic AI
Workflow Execution	Static workflows	Dynamic multi-step workflows
Adaptability	Limited	Highly adaptive
Autonomy & Initiative	Reactive	Proactive
Context Awareness	Narrow scope	Holistic view
Collaboration	Minimal	Human and agent collaboration
Learning & Improvement	Manual updates	Continuous learning

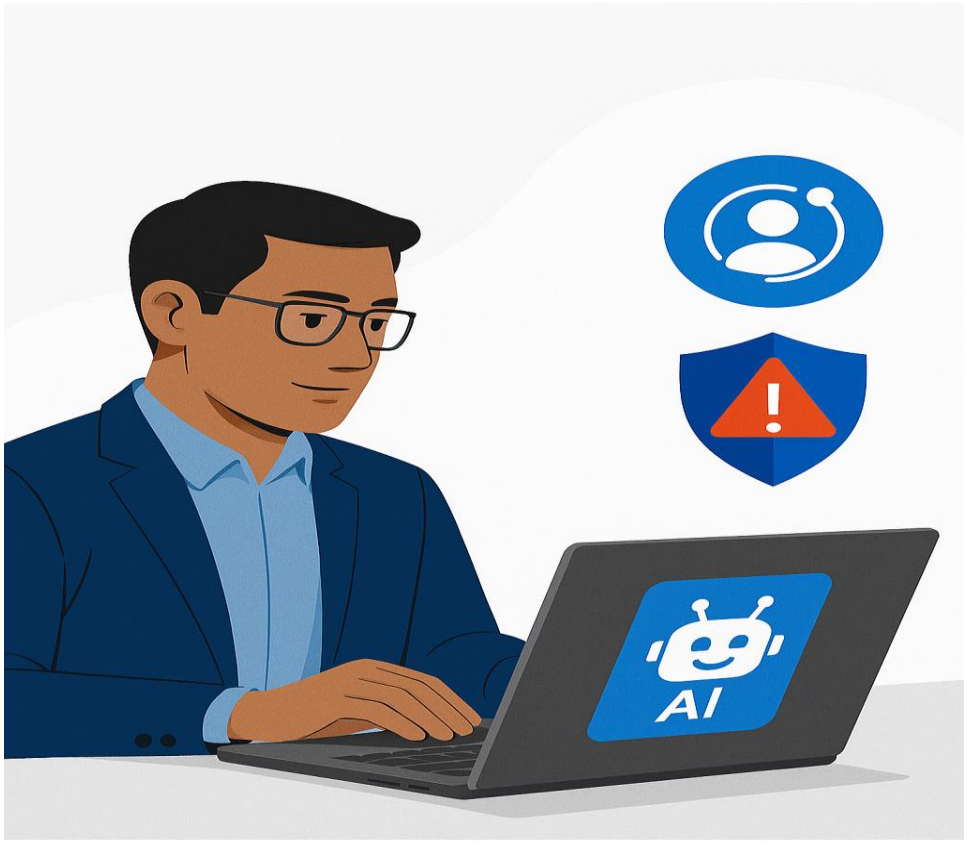

Agentiic AI

Agentic AI: A New Approach to Vetting

Agentic AI refers to AI systems that act as autonomous agents, capable of planning, executing, and adapting multi-step tasks independently in dynamic environments. Unlike traditional AI, which typically follows a fixed, rule-based algorithm or handles single tasks in isolation, an agentic AI can manage an entire workflow on its own.

Agentic AI Solution: Architecture & Technology

The core of the transformation was the adoption of Agentic AI—autonomous agents capable of handling complex, multi-step verification tasks. Key features include:

- Autonomous AI agents for multi-step verification
- Automated data gathering, partner history analysis, and decision-making
- Fraud risk detection through document and database analysis
- Transformation of linear workflows into collaborative, AI-driven systems
- Integration and use of tools like SharePoint Agents, Microsoft Teams, Copilot Studio, Azure AI Foundry, Microsoft Agent Framework, and GitHub Copilot

The high-level architecture supports scalable, secure, and efficient vetting, with robust governance and compliance controls.

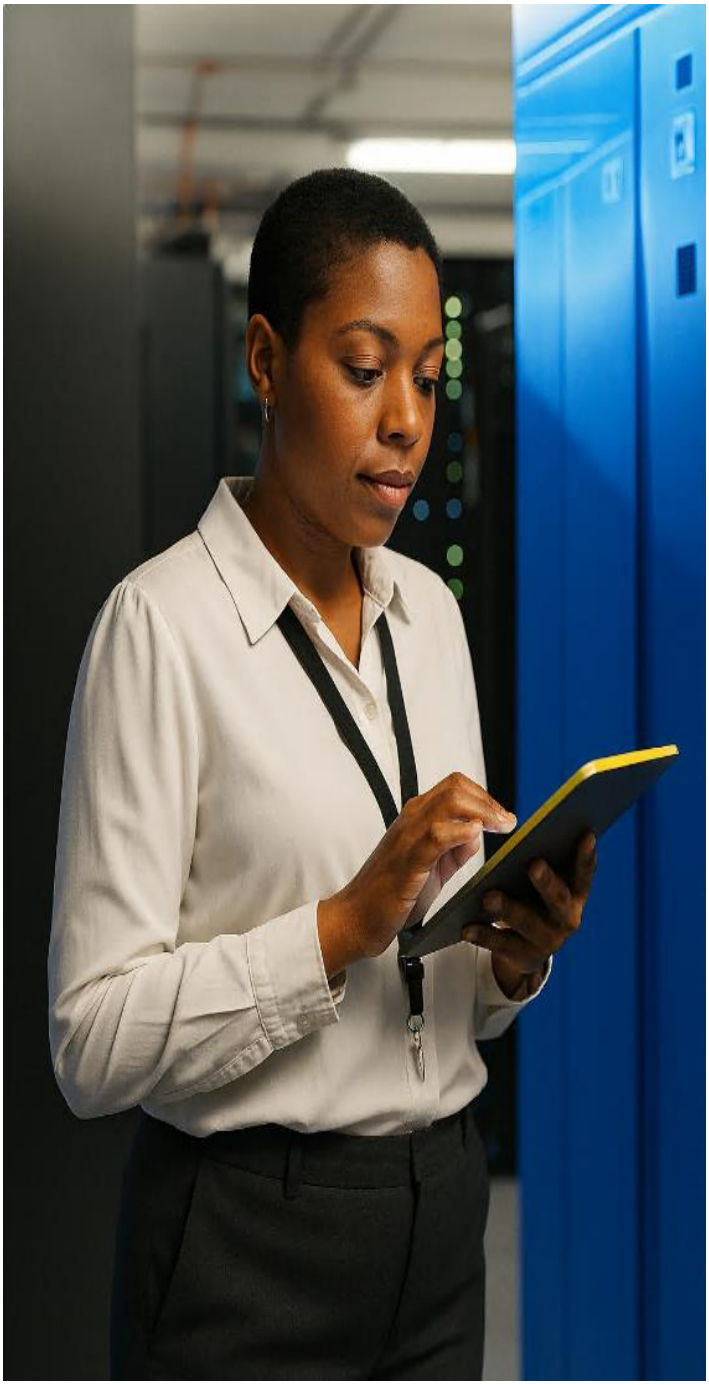

Idea Storming – Art of the Possible (Think AI!)

Before designing the Agentic AI solution, a collaborative idea-storming session was conducted to envision the full potential of AI in transforming the vetting process. Key questions explored included:

Which review steps can AI agents handle autonomously?

- Identifying repetitive, rules-based, or data-intensive tasks that can be fully automated by AI agents, freeing human reviewers for more complex cases.

How can AI agents surface key insights for reviewers?

- Leveraging AI to analyze large datasets and highlight critical information, anomalies, or trends that might otherwise be missed.

Where can AI agents flag risks before humans spot them?

- Using predictive analytics and pattern recognition to proactively identify potential fraud or compliance risks earlier in the process.

How do AI agents outperform pattern recognition?

- Applying advanced machine learning models to detect subtle, non-obvious patterns and correlations that are difficult for humans to discern.

Can AI agents streamline reviewer collaboration?

- Facilitating seamless information sharing and task coordination among reviewers through intelligent workflow automation.

How might AI agents reduce manual workload and lag?

- Automating routine checks and data gathering, significantly reducing processing times and reviewer fatigue.

What blind spots can AI agents help uncover?

- Continuously monitoring emerging risks, new fraud tactics, or overlooked compliance gaps, ensuring a more robust and adaptive vetting process.

These questions guided the design and implementation of the Agentic AI solution, ensuring that the transformation addressed both current pain points and future opportunities.

Technology & Architecture

Core Framework Used

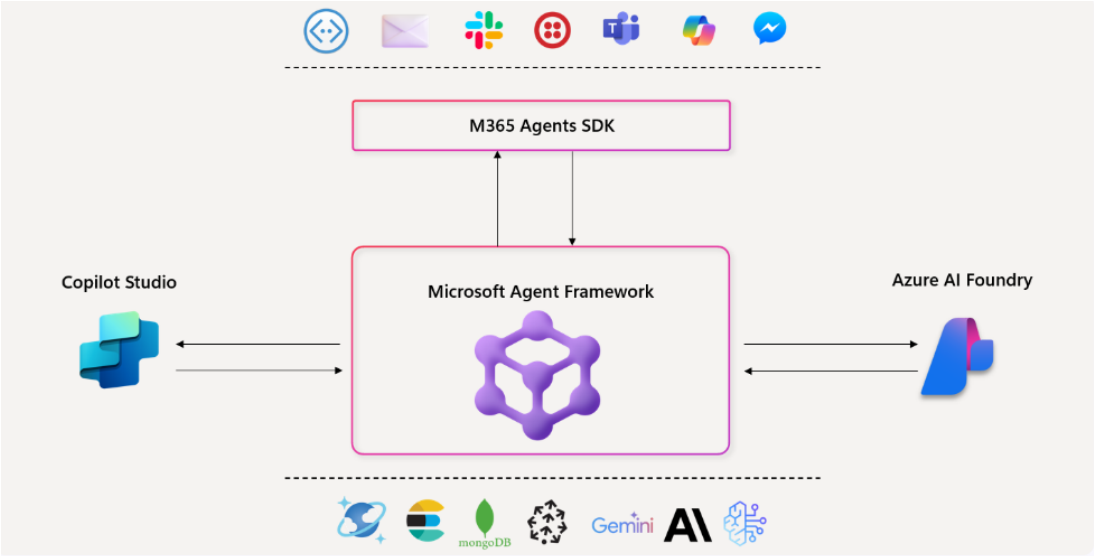

High Level Architecture

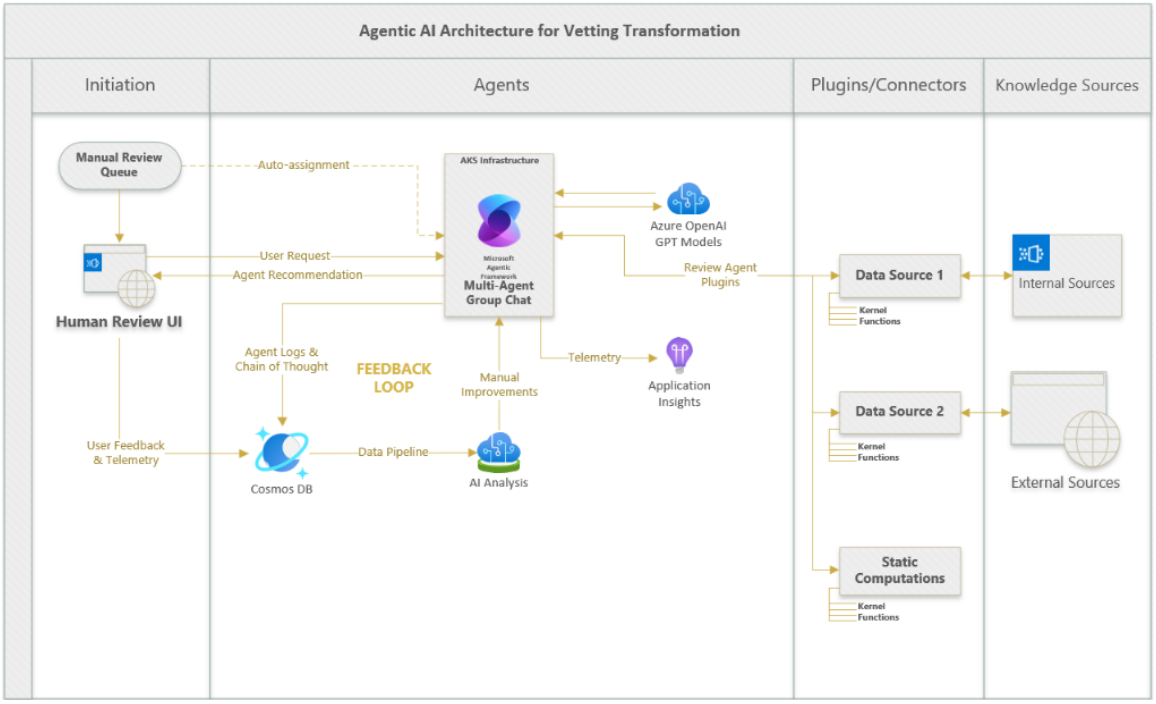

Key pointers about this architecture:

- The heart of the design is a multi-Agent system running on Microsoft Agent framework deployed on Azure Kubernetes services.
- Microsoft Agent Framework provides Multiple External Integrations using Plugins. Any Internal or External signals are injected using plugins.
- There is integration with Azure services to store Telemetry, Logs and Chain of thoughts.
- Logs & Chain of thoughts are later reused for better decision-making using AI Analysis service.
- It has inbuilt connectivity with Azure AI foundry and Copilot Studio.

Agentic AI Solution: Architecture & Technology

The transformation delivered substantial, quantifiable benefits:

- **SLA Improvement:** Reduced onboarding from days to minutes
- **Task Efficiency:** Individual task times dropped from days to minutes.
- **Analyst Time Saved:** Over 150K estimated hours saved for a year.
- **Always-On Monitoring:** Consistent, fatigue-free oversight
- **Cost Reduction:** Significant operational savings and development efficiency

Overview & Capabilities of new Agentic Solution

1) New system uses autonomous AI agents to handle complex multi-step verification tasks previously done by humans.

2) Agents automatically gather related info, analyze partner history, and decide verification outcomes or escalate cases.

3) AI agents comb through documents and databases to detect potential fraud risks during verification and due diligence.

4) New System transforms linear workflows into collaborative systems where AI handles data collection and initial analysis.

Most popular Agentic AI Frameworks offered by Microsoft for different Audience & Skill levels

Tool	Users	Skill Level	Features
SharePoint Agents	Business users	No coding	Document-centric automation
Microsoft Teams & M365 Copilot	General professionals	No coding	Contextual AI assistance
Copilot Studio	Power users	No code	Custom AI workflows
Azure AI Foundry	AI engineers	Low code	Scalable ML services
Microsoft Agent Framework	Developers	Pro-code	Language understanding SDK
GitHub Copilot	Developers	Pro-code	AI-powered coding

Impact and Measurable Results of Transformation

Metric	Before	After	Improvement
Onboarding SLA	Days	Minutes	Accelerated
Task Time	>20 min	<1 min	95% reduction
Analyst Workload	High	Reduced	150k+ hours saved
Availability	Business hours	24/7	Continuous operation
Cost	High	Low	Significant savings

Best Practices followed during this Transformation Journey

Key lessons from this transformation journey include:

- Promote adoption early by educating leaders and frontline teams
- Inject AI throughout the process for maximum impact
- Enhance skills and pilot initiatives to build trust
- Engage users by positioning AI as a helpful partner
- Integrate responsible AI and security from the outset
- Monitor and refine continuously using telemetry and feedback

Responsible AI Principles followed during the process

- **Fairness** AI systems should treat all people fairly.
- **Reliability and safety of** AI systems should be performed reliably and safely.
- **Privacy and Security** AI systems should be secure and respectful for privacy.
- **Inclusive** AI systems should empower everyone and engage all people, regardless of their background.
- **Transparency of** AI systems should be understandable.
- **People** should be accountable for AI systems.

Broader Applicability

The Agentic AI framework developed in this project is designed to address the challenges of complex, multi-step processes that traditionally require significant manual effort. Its architecture enables autonomous agents to plan, execute, and adapt tasks across dynamic environments, making it highly versatile. While the initial focus was on streamlining entity vetting, the underlying principles—such as intelligent automation, continuous learning, and collaborative decision-making—can be applied to a wide range of operational scenarios. By reducing reliance on manual intervention and enhancing process efficiency, Agentic AI helps organizations achieve faster, more accurate outcomes.

Beyond entity vetting, this framework holds promise for transforming other domains where manual workflows are prevalent. For example, compliance management, fraud detection, and customer support all involve intricate procedures that benefit from automation and adaptive intelligence. By leveraging Agentic AI, organizations can improve risk mitigation, ensure regulatory adherence, and deliver better customer experiences. The flexibility and scalability of the framework make it an asset for any enterprise seeking to modernize its operations and drive continuous improvement across multiple business functions.

Conclusion

The vetting transformation at Microsoft stands as a pioneering example of how Agentic AI can be harnessed to drive meaningful change in enterprise operations. By integrating autonomous intelligent agents into the partner vetting process, Microsoft achieved remarkable improvements in speed, accuracy, and efficiency—reducing onboarding times from days to minutes and saving over 150,000 analyst hours in a single month. This initiative not only addressed legacy challenges but also sets a new benchmark for fraud detection and compliance, demonstrating the tangible business impact that advanced AI solutions can deliver. The project's success was rooted in a holistic approach that combined innovative technology with a clear vision, disciplined execution, and a commitment to continuous improvement.

Beyond the immediate operational gains, this transformation highlights Microsoft's thought leadership in AI-driven process innovation. The Agentic AI vetting framework serves as a trailblazing model for other organizations seeking to modernize complex workflows and elevate standards of excellence. It offers valuable lessons on the importance of aligning technology with business objectives, fostering cross-functional collaboration, and maintaining agility in the face of evolving risks. As enterprises continue to navigate the challenges of digital transformation, Microsoft's experience provides a roadmap for leveraging AI not just for efficiency, but for sustained excellence and industry leadership.