



Frontier Finance Brief

Security and Responsible AI Considerations for Finance Professionals

Author:


Lina Rebouh

Key Contributors:

Duane Kamihara, Mike Tokic, MJ Aldred, Greg Johnson,
Russel Hercules

Tags:

Finance Agents, Copilot Studio, Generative AI, Responsible
AI, Security



Overview

Microsoft's Secure Future Initiative (SFI) and Office of Responsible AI (RAI) offer guidelines for finance professionals to enhance artificial intelligence (AI) capabilities while maintaining system security and reliability. This brief will cover SFI and RAI applications relevant to finance professionals, as well as explore finance use cases using Microsoft Finance Agents and Copilot Studio.

Why is it important

The integration of AI within the financial sector is advancing rapidly, presenting both significant opportunities and challenges due to the sensitive nature of financial data. As AI applications expand within the industry, addressing security concerns becomes imperative. Successfully navigating AI projects involves managing these risks and ensuring the deployment of AI is secure, ethical, and responsible. This article will explore how finance professionals can navigate this journey in a secure and responsible manner.

Microsoft Secure Future Initiative (SFI)

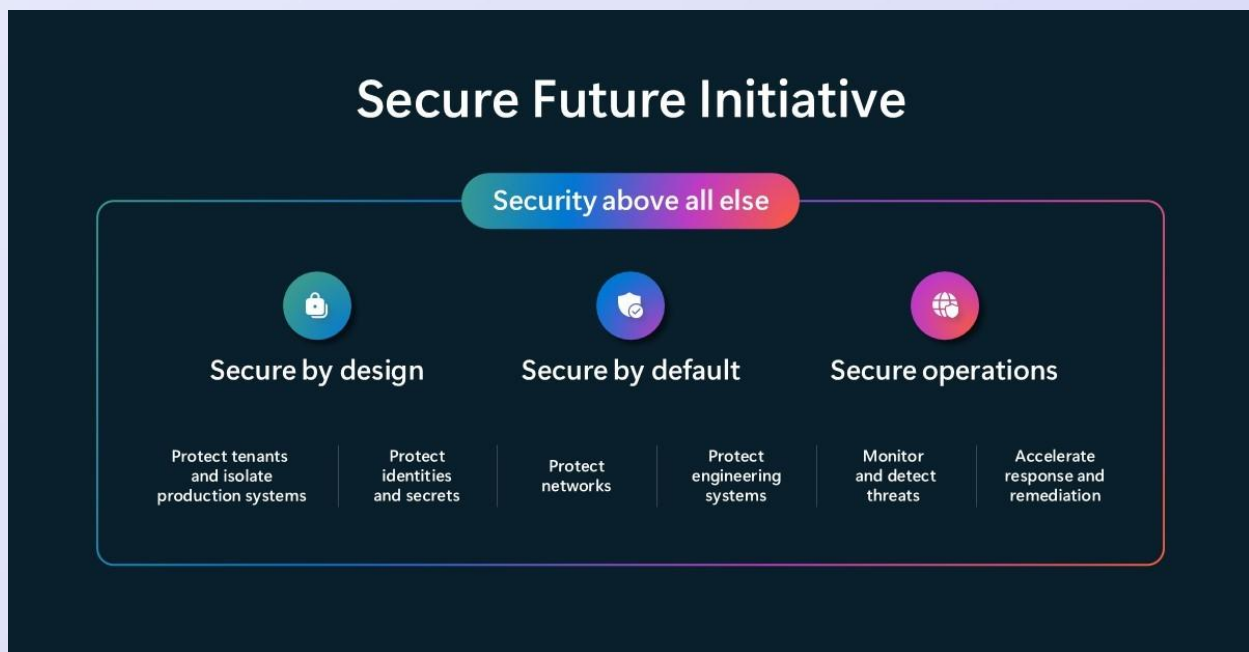
For finance professionals, addressing security topics may initially seem daunting. While it may not be the primary concern of a typical finance role, understanding key security concepts and leveraging available resources is essential for creating secure and compliant AI solutions to support various finance scenarios. At Microsoft, this is achieved through the Secure Future Initiative (SFI) framework. This comprehensive framework aims to continuously enhance security measures across Microsoft's products and services, ensuring they meet the highest standards to protect against evolving cyber threats.



Launched in November 2023, the Secure Future Initiative is based on three core principles: Secure by design, secure by default, and secure operations. The principles are implemented across six pillars: protect tenants and isolate production systems, protect identities and secrets, protect networks, protect engineering systems, monitor and detect threats, accelerate response and remediation. [Learn more](#)

“AI transformation requires Security Transformation.”

Charlie Bell, EVP of Microsoft Security, Microsoft Ignite Keynote, 2024





Office of Responsible AI (RAI)

While security addresses protecting financial data, Finance professionals must also consider building AI solutions ethically. At Microsoft, this is achieved through the Responsible AI initiative. This program, launched in 2018, ensures AI systems are designed and used responsibly, guided by six core principles: Fairness, reliability and safety, privacy and security, inclusiveness, transparency and accountability. [Learn more](#)



Fairness

AI systems should treat all people fairly.



Reliability and safety

AI systems should perform reliably and safely.



Privacy and security

AI systems should be secure and respect privacy.



Inclusiveness

AI systems should empower everyone and engage people.



Transparency

AI systems should be understandable.



Accountability

People should be accountable for AI systems.

How do we apply these principles in finance? Microsoft encourages its employees to [champion Responsible AI](#). This approach ensures the ethical implementation of AI throughout the company. Finance professionals

"Microsoft runs on trust, and our success depends on earning and maintaining it. We have a unique opportunity and responsibility to build the most secure and trusted platform that the world innovates upon."

Satya Nadella, CEO of Microsoft

collaborate with the Office of RAI to create governance frameworks, define roles, and promote a culture of readiness for RAI.

Within this framework, finance professionals learn about a risk management approach that covers the entire AI development lifecycle. This includes impact assessments and [red teaming](#) to identify potential risks and ensure additional oversight and support for teams developing high-risk models and systems through governance processes.

The RAI finance champion community considers various finance AI scenarios to understand the unique sensitivities of each use case. The collaborative efforts between the Office of RAI and its champion community are essential for maintaining fairness and transparency. Let's now dive into an example of a use case and see how we can put security and RAI principles into practice.

"Our Office of Responsible AI prioritizes keeping humans not just in the loop, but at the center of AI systems"

Natasha Crampton, Vice President and Chief Responsible AI Officer, [AI Impact | Microsoft CSR](#)

Finance Use Case: Account reconciliation with Microsoft Finance Agents in Excel

Consider a scenario where a Finance professional is tasked with the monthly reconciliation of accounts, which involves the below manual tasks:

- Manual analysis of Finance mapping keys
- Reconciliation analysis and listing of matched, potentially matched and unmatched transactions



- Reconciliation report generation with summary and insights of results found
- Transaction troubleshooting for any unreconciled line items.

By leveraging Generative AI tools such as Microsoft's Financial Reconciliation Agent, these tasks can be automated, leading to efficiencies and time-savings in the monthly account reconciliation process, allowing Finance professionals more time to investigate the potential data discrepancies.

Finance Agents security considerations

Here are some common security questions Finance professionals might have before using Finance Agents to automate tasks within their processes.

In what ways do Finance Agents demonstrate their commitments to security?

Finance Agents adhere to Microsoft's SFI, ensuring security by default, design, and operations.

What sets of data do Finance agents access to?

The Finance Agents data layer consists of three distinct sets of data. Microsoft Graph (MS Graph) data (end user's existing Microsoft 365/Office Data), ERP Data (stored in the existing ERP system), Non-ERP Data (in customer's Power Platform Dataverse), and Finance Agents data (agents generated insights data [Architecture overview](#)).

How is MS Graph data, and the other sources of data utilized? Will Finance Agents assist in training large language models (LLMs)?

Prompts, responses and data accessed through your Microsoft Graph aren't used to train LLMs, including those used by Microsoft 365 Copilot ([Data, Privacy, and Security for Microsoft 365 Copilot](#)). Finance Agents always respect the data privacy, security, retention and compliance boundaries of the underlying data store for data at rest ([Architecture overview](#)).

How does access governance work?

As a prerequisite, Power Platforms administrators must first enable the Financial Reconciliation Agent within a given Dataverse environment. The implication of this deployment approach means organizations can easily control which users are going to have access to a given agent. Access is granted based on the end user's authentication context, aligning with the permissions assigned to them. This means that if a user has specific access within a financial system, they will see the same data results in Finance Agents (Architecture overview).

What are the interactions with open-source models?

There are no interactions between Finance Agent data (MS Graph, ERP or non-ERP Data Integration layer), and Open AI Service, as per the architecture ([Architecture overview](#)).

Finance Agents RAI considerations

Finance Agents adhere to Microsoft's RAI principles. Below are questions finance professionals might have when using Finance Agents and how they apply these principles.

What operational factors allow for effective and responsible use of Finance Agents in Excel?



Users are required to follow the specific data requirements when formatting data into tables. Once done, they receive suggestions for reconciliation vectors, and can review, accept, or override Finance Agents suggestions. **The product is designed to put the human at the center and empower the user to use their own judgement and scrutiny prior to consuming the results** ([Architecture overview](#)).


“An AI system includes not only the technology, but also the people who use it, the people affected by it, and the environment in which it’s deployed”

[Responsible AI FAQs for Finance Agent](#)

What is an example of Accountability in action in Financial reconciliation agent in Excel?

After a reconciliation report is generated, users have the ability to **“Append summary to report”**, which generates a textbox summary at the bottom of the tables of matched and unmatched transactions.

Users can edit the summary as applicable by overwriting the results suggested by Finance Agents.



Reconciliation report summary

- Here is an example of modification I can make to the Report summary. The reconciliation between the T BankStatementTable) resulted in 13 set of transactions being perfectly matched. These transactions had combination of Date and Expenditure category with Transaction date and Transaction category. Here is a


Finance Use Case: Account reconciliation agent


Building upon the previous use case, let's create a new financial reconciliation agent that can autonomously start financial reconciliations on the user's behalf and notify them when the process is completed while putting both security and RAI principles into practice. Tasks the agent can help with include:


- Automate the process of collecting data from specified enterprise financial data source
- Automate starting the reconciliation process based on predefined trigger
- Automatically notify specific users when the reconciliation is completed

These steps can be completed directly in Microsoft Finance agents in Excel or automated based on the configured trigger.

Agents can work on your behalf, performing tasks autonomously based on prescribed instructions and selected connectors. In Excel, in the Finance agents' sidecar, the user can



 Agent name

 Data source

SharePoint site url *

Library *

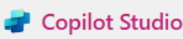
Folder hierarchy *






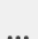
Folder path *

select "Agents." With a few clicks, the user can fill out data source location and add trigger, which can either be an event, or a schedule, and add recipients of the report.

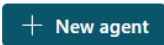
Recipient(s) *

The selected recipients will be notified via Outlook email.




 Home
 Create
 Agents
 Flows
 Tools
 ...


Create




Install a managed agent (preview) ⓘ




Awards and Recognition
 Managed agent



Employee Self-Service
 Managed agent




Financial Reconciliation (Preview)
 Managed agent



Leave Management
 Managed agent


Start with an agent template



Safe Travels
 Agent template

Provides answers to common travel questions and related health and safety guidelines.

Government



Website Q&A
 Agent template

Instantly answer user questions using the content of your website or other knowledge.

AI + Machine Learning

Customer Service

One of the beneficial features is that the autonomous agent can be configured to operate differently from the one in Excel, within Copilot Studio. For example, finance users may want agents to initiate the reconciliation



process when three specific conditions are met, gather data from both an ERP system and another source, execute the reconciliation, and notify users via Teams instead of email. All these functionalities can be achieved in Copilot Studio, without compromising Security or Responsible AI principles.

Copilot Studio security considerations

Copilot Studio adheres to Microsoft's SFI, ensuring security by default, design, and operations. Prior to starting the adoption and building an agent in Copilot Studio, Finance professionals might want to engage with their compliance team, to address security questions.

The Copilot Studio team has compiled a list of [Security FAQs](#) to help Finance professionals with some common questions:

- How do tenants work with Copilot Studio?
- How do connectors with Microsoft Power Platform work?
- How can Power Platform admins within your organization control generative AI capabilities in Copilot studio?

Copilot Studio RAI considerations

RAI considerations in Copilot Studio are available for Finance professionals to review while building custom agents: [FAQ for generative orchestration](#).

Finance professionals at Microsoft use the Responsible AI Impact Assessment, which is now public. Use the Microsoft Impact Assessment template ([Microsoft-RAI-Impact-Assessment-Template](#)), or your organization's equivalent to align with IT, Security, and Finance teams for responsible AI solutions.

The template is structured to evaluate the effects of AI systems on individuals, organizations, and society. It comprises sections for:

1. System information
2. Intended Uses
3. Adverse Impact
4. Data Requirements
5. Summary of impact

Conclusion

Although there are infinite use cases for leveraging AI to generate efficiencies within the finance function, there is a finite path to ensure that the developed AI solutions are secure, compliant, and responsible.

"As we design, build, and release AI products, six values – transparency, accountability, fairness, inclusiveness, reliability and safety, and privacy and security – remain our foundation and guide our work every day.

Satya Nadella, CEO of Microsoft

References

- Microsoft Secure Future Initiative (SFI)
- Secure Future Initiative (SFI) April 2025 progress report
- Expanding Microsoft's Secure Future Initiative (SFI)
- Prioritizing security above all else - The Official Microsoft Blog
- Responsible AI at Microsoft



- Microsoft Responsible AI Standard v2
- Architecture overview
- Empowering responsible AI practices
- Responsible AI Principles and Approach | Microsoft AI
- Innovating in line with the European Union's AI Act
- Microsoft AI Red Team
- The building blocks of Microsoft's responsible AI program