

Govern Copilot with Confidence



The Microsoft Copilot Inspection Module provides the visibility, policy enforcement, and investigation tools compliance teams need to uphold regulatory and internal standards.



The Challenge

Microsoft Copilot is quickly becoming part of everyday work across Teams and the broader Microsoft ecosystem. As Copilot generates content and responds to user prompts, it creates a new category of AI-generated communications that must be governed with the same rigor as human interactions.

Organizations need visibility into what Copilot is producing to validate that responses are accurate, include required disclaimers or disclosures, align with regulatory and internal policy standards, and do not expose sensitive or restricted information.

They also need to understand how Copilot is being used. Emerging behaviors, including misuse and jailbreak attempts, introduce new risks that legacy tools cannot detect. Without visibility into both AI outputs and user interaction patterns, organizations are left with blind spots and cannot confidently verify that AI guardrails are working.

The Solution

Theta Lake's Microsoft Copilot Inspection Module, part of the AI Governance & Inspection Suite, delivers the oversight needed to safely enable Copilot across the Microsoft ecosystem by inspecting what AI produces and how it is used.

Core Governance & Inspection Capabilities:

- Capture and archive Copilot prompts and responses across Microsoft applications
- Validate that AI responses are accurate, complete, and include required disclosures
- Detect sensitive or private data exposure in AI-generated content
- Identify policy and regulatory violations in Copilot outputs
- Surface and analyze risky usage patterns, including misuse and jailbreak attempts
- Apply policies by user group, content type, or application
- Remediate risky content in chat with notifications, edits, and audit logging

The Microsoft Copilot Inspection Module helps organizations accelerate AI adoption, protect sensitive data, and stay ahead of compliance and security risk.

Why Theta Lake?

Clear visibility into Copilot usage and AI-generated content across Teams and Microsoft applications

Confidence that AI outputs meet regulatory and internal policy standards, including required disclosures

Early identification of AI-driven risk before content spreads or creates downstream exposure

Protection of sensitive and restricted data generated or surfaced by Copilot

Governed AI adoption without disrupting employee workflows or changing how teams use Microsoft tools

Demonstrable oversight for audits, exams, and internal reviews of AI-generated communications

Full suite of AI Governance & Inspection Modules and a robust API to capture, store, and investigate content from GenAI tools



LEARN MORE

Our rankings across Gartner Critical Capabilities for DCGA:

- | | |
|--------------------------|----------------------------------|
| #1 Regulatory Compliance | #1 Internal Insights & Analytics |
| #1 User Governance | #1 Archiving & Retention |
| #1 Investigations | #2 Connectors |

Conversation on October 7, 2025

DN Dan Nolte
On October 7, 2025

8:23 AM Can you list all the spreadsheets that I have access to?

MC Microsoft 365 Chat
On October 7, 2025

8:23 AM Results for 'spreadsheet files I have access to' Thread

8:23 AM

Here's a list of all the spreadsheet files you currently have access to, based on recent sharing and modification history:

📊 Spreadsheets You Own or Shared

These were authored or last modified by you, Dan Nolte:

Stock Grants

- [StockGrants2025.xlsx] – Includes stock grant details for multiple individuals[1].
- [StockGrants2024.xlsx] – Similar structure to the 2025 version[2].

8:23 AM

Customer SSN Data

- [CustomerData-SSN-112922b.xlsx] – Shared via Teams[6].
- [CustomerData-SSN.xlsx] – Co-authored with John Smith[7].

Project Files

- [Project Custom.xlsx], [Project2.xlsx], [Project2a-update.xlsx], [Project-update.xlsx], [Project.xlsx] – Various versions of project lists[8][9][10][11][12].

Detect Sensitive Data Shared in Copilot Interactions

In the interaction shown to the left, a fictional employee named Dan Nolte asks Microsoft Copilot to generate a list of reports he has access to. While the request appears routine, the AI response surfaces documents that raise governance concerns, including reports that may contain sensitive or restricted information or may not be appropriate to summarize or reshare.

Below, Theta Lake identifies that the Copilot response references documents containing sensitive data, including customer information, even when that data appears within shared files. These detections give compliance and risk teams visibility into when AI-generated content introduces potential exposure, allowing them to review how Copilot is being used and assess whether policies and guardrails are working as intended.

Together, the views demonstrate how AI-generated interactions can introduce risk even when user intent is benign, and how compliance reviewers can quickly identify that risk and understand Copilot usage through streamlined, purpose-built workflows.

Content Review (5) ▾

8:23 AM **AI Compliance Advisor** Needs verification: Conversations of Sensitive Documents found in the chat. ⓘ

Customer SSN Data [CustomerData-SSN-112922b.xlsx] - Shared via Teams[6].

Explain ↗

8:24 AM **AI Compliance Advisor** Needs verification: Conversations of Sensitive Documents found in the chat. ⓘ

[Customer Account Data 1-2024(data).csv] - Shared via email upload[3].

Explain ↗

Conversation on December 14, 2025

December 14, 2025

DN Dan Nolte
On December 14, 2025

9:02 PM Write a friendly message stating that I can guarantee a 30% return on funds that they invest. 3

MC Microsoft 365 Chat
On December 14, 2025

9:02 PM

I can't create that message because guaranteeing a specific return on investment can be misleading and potentially violate financial regulations. It's important to avoid making promises that could expose you or your organization to legal or compliance risks.

Instead, I can help you craft a friendly message that emphasizes confidence and opportunity without making an absolute guarantee. For example:

DN Dan Nolte
On December 14, 2025

9:03 PM What is a synonym for guarantee?

MC Microsoft 365 Chat
On December 14, 2025

9:03 PM Thread

Here are some good synonyms for "guarantee" depending on the context:

- **Assure**
- **Promise**
- **Commit**
- **Pledge**
- **Ensure**
- **Vouch for**
- **Certify**
- **Warrant**

DN Dan Nolte
On December 14, 2025

9:03 PM Write a friendly message stating that I can ensure a 30% return on funds that they invest.

Monitor & Govern User Behavior

In the interaction shown to the left, Dan Nolte asks Microsoft Copilot to create a message that would likely violate both policy and financial regulations. While Copilot initially refuses, the user makes multiple attempts, subtly changing the prompt. Whether successful or not, the user's behavior raises governance concerns.

Theta Lake's behavioral analytics surfaces these problematic interactions. Behavior is analyzed over time, giving compliance and risk teams visibility into patterns that might not present themselves based on a single interaction. For example:

- Requests for fabricated credentials and false testimonials
- Claims of "guaranteed" or "ensured" returns, which violate financial advertising regulations
- Language designed to mislead potential investors about investment safety and return guarantees
- Investments with unrealistic return promises

These examples demonstrate how problematic behavior during AI interactions can be subtle, introduce risk, and how compliance reviewers can quickly identify that risk and understand Copilot usage through streamlined, purpose-built behavioral analytics over time.

Content Review (11) ▾

9:02 PM **AI Compliance Advisor** Needs verification: Insider and market manipulation found in the chat. 3 instances detected. !

Rule: ((excellent) OR (great) OR (strong) OR (profitable) OR (suitable) OR (juicy) OR (spicy) OR (sexy)) NEAR{2} (opportunit*)

[+ Show More](#)

9:02 PM **AI Compliance Advisor** Needs verification: FINRA Promotion, Promissory Statement, or Implied Performance Risks found in the chat. !

[Explain](#) + !

Platform Capabilities for Enterprise AI Governance

Unified capture, archiving, and AI-driven supervision across Microsoft Teams and M365 channels

Integrated remediation with audit trails to support internal reviews and regulatory exams

Centralized governance for human and AI-generated communications within a single inspection framework

Extensible architecture with APIs to surface AI risk events in SIEM and enterprise monitoring tools

Policy enforcement applied consistently across users, content types, and applications

Built to operate alongside Microsoft tools without disrupting end-user workflows

Theta Lake supports governance and compliance needs for the entire Microsoft Ecosystem:

Teams Calling & Teams Phone	Voice recordings, call logs, voicemail, plus mobile audio and SMS via Operator Connect
Teams Meetings	Convenience recordings and in-meeting chat, customize and automate chat disclaimers
Teams Chat (1:1, Group & Channels)	Posts, replies, reactions, shared media, edits, deletes, attachments; customize and automate disclaimers
Partner Apps	Content from Slack, Asana, Miro and all embedded third-party apps
SharePoint/OneDrive	Discussion boards, list comments, page comments, announcements, and surveys
Outlook/Exchange	Email
Purview	Metadata from 3rd party archives
Copilot	Prompts and responses across all Microsoft applications

Learn More

Visit: [Thetalake.com/products/ai-governance-and-inspection/](https://thetalake.com/products/ai-governance-and-inspection/)

