When private data becomes public knowledge

Not all information is intended for everyone

Copilot makes it easier to find and summarise information across Microsoft 365. It works in line with the security policies and controls in place, so it's important that your environment is set up correctly.

Introducing AI into an environment built on inherited permissions, outdated access, and inconsistent policies increases the likelihood of data leaks or exposure.

Without clear boundaries, Copilot can surface sensitive data to the wrong people, such as:

- Unredacted transcripts from private Teams meetings
- Salary information, bonus structures or employee personal details
- Commercial or financial data from supplier negotiations

Copilot doesn't grant access, but it does accelerate the access that already exists (even if it's incorrect).

Microsoft 365

What this is

Access to sensitive data is granted by default

Your HR department stores everything from policy documents to employment contracts in a dedicated SharePoint site. Permissions haven't been reviewed for years, and access was inherited from a wider group. Copilot directs employees to files containing sensitive information about their colleagues simply because they are already part of that group.

The accidental sharing of trade secrets

Someone asks Copilot to summarise internal discussions to send to a supplier. At first glance the response reads like a simple recap. But the summary includes pricing assumptions, negotiation points, and internal commentary pulled from the underlying emails and Teams chats. Your commercial strategy is now sitting in your suppliers inbox (and your bargaining power is on the floor).

Malicious instructions hidden in emails

A team member asks Copilot to summarise an email thread. Hidden inside one of the messages are instructions designed to influence the output from AI. The text attempts to push the system to reveal additional information or bypass normal safeguards. This type of cyber threat is known as prompt injection, where malicious instructions are embedded in documents or emails to manipulate AI responses.

What leaders can do about it

Control can weaken as teams change and cloud environments grow. It's important to have a clear understanding of who can access what before bringing AI in.

Here are three priorities that leaders should focus on:

1. Validate governance rather than assuming it exists

Review policies, processes and controls and gather evidence that they are actually being applied. If a control cannot be demonstrated in practice, it should not be assumed to be working.

2. Take stock of who has access to sensitive information

Establish who can access HR files, financial data, board documents and supplier negotiations. In many organisations, permissions have expanded gradually over time and no longer reflect current roles.

3. Evaluate AI readiness before deployment

Conduct an AI readiness assessment to identify exposure risks and governance gaps before AI tools are introduced at scale. Some organisations carry this out internally, while others engage a specialist third party to provide independent insight and technical expertise.

Book an AI Readiness Assessment

Understand what your environment could reveal before introducing AI into everyday work. Identify existing data exposure risk and improve governance ahead of AI adoption.

Book your readiness assessment with 848 Group today.

Book an assessment