proarch

Managing AI Risks: How to Protect Sensitive Information with Microsoft Purview

Presented by: David Trum & Todd Brink

MEET OUR PRESENTERS

David Trum

Senior Solutions Strategist
ProArch

Todd Brink

Information Security Senior Consultant
ProArch

Agenda

- Why Monitor AI?
- What is the downside?
- Sanctioned vs. Unsanctioned GenAI
- Why perform Access Audits?
- Sensitivity Labels
- Limiting Unsanctioned GenAI on company devices
- Discovery
- Control

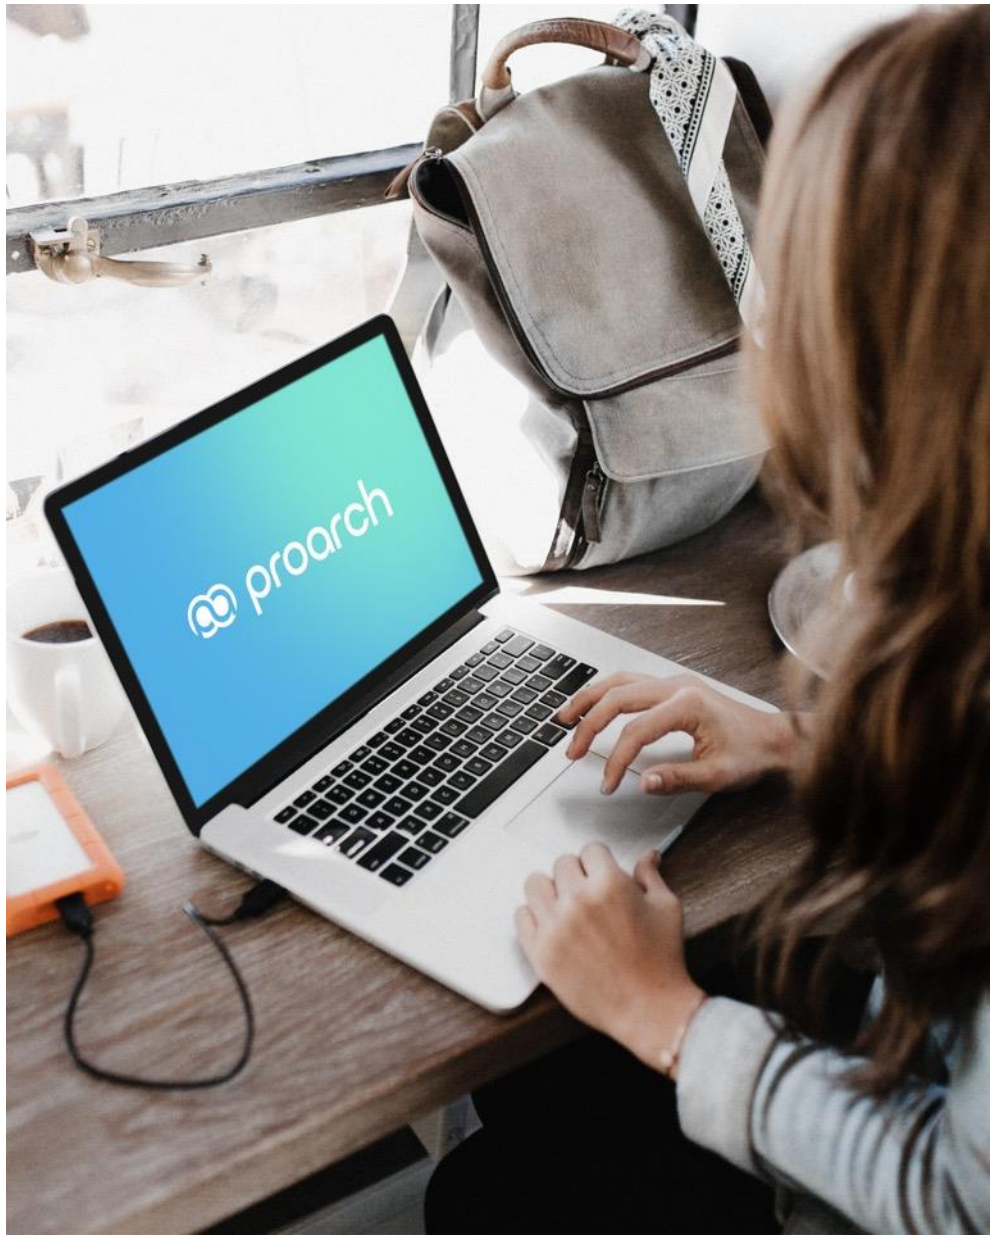

Generative AI offers a new method of searching data, making anything within your scope accessible with just a prompt.

It is an amazing tool to increase productivity and leverage intellectual capital.

This is an addressable risk.

"Security by Obscurity" is gone.

Just ask the right question.

GenAI and Copilot are powerful tools for enhancing productivity, but they pose new risks due to the ways people might input data.

The bigger the organization, the bigger the risk because the more users, rules, and data there are.

AI Gains Bring Data Challenges

Data Oversharing

- Lack of labeling policies or access controls can lead to unauthorized access to sensitive data via AI apps.
- Users may view or edit sensitive data without proper authorization.
- AI apps can expose data if proper access controls are not in place.

Data Leakage

- Users may leak sensitive data to unsanctioned AI apps.
- Sanctioned apps can also pose risks if AI responses don't inherit data protection controls.

Noncompliant Usage

- Users may generate high-risk content with AI apps.
- AI apps can create content that violates ethics standards.
- AI-generated documents may conceal illegal activities like insider trading or money laundering.

Consumer vs. Enterprise GenAI

Governance challenges increase with many users accessing GenAI with their files and data.

Managing a single consumer-grade GenAI focused on public information is much simpler.

The complexity of facilitating multiple users with varied data is significantly higher.

Finding a sweet spot between data protection and productivity can paralyze the system

Microsoft Purview: comprehensive data security

- Gain visibility into data across your organization
- Safeguard and manage sensitive data across its lifecycle, wherever it lives
- Manage critical data risks and regulatory requirements

proarch

Sanctioned vs. Unsanctioned

Sanctioned vs. Unsanctioned

SANCTIONED AI

The organization retains some control.

Traditional access controls are respected – users just have a new way of getting at the information.

Mitigate this using:

1. Periodic access audits
2. Clean up old sharing links
3. Sensitivity labels

UNSANCTIONED AI

It's a free for all.

Endpoint agents like Defender Endpoint DLP and edge content sensitivity tools, monitor and block apps from being used.

Sanctioned AI Example: Copilot for Microsoft 365

Copilot for Microsoft 365 was asked a broad prompt:

What should I focus on today?

Copilot responded:

Focus on reviewing the Performance Improvement Plan (PIP) for John Smith in Sales.

Why did this happen?

Faulty permissions in SharePoint that could have been avoided with Sensitivity Labels.

Sanctioned AI Example:

How to prevent prompts from surfacing sensitive information

Copilot for Microsoft 365 adheres to the same file and item permissions that are set within your organization.

- Explicit Access Controls
- Sharing Links

Emails & Attachments	Calendar Entries
SharePoint	Shared Calendars
OneDrive	To Do & Planner
Teams Chats	Yammer
Teams Channels	OneNote in SharePoint
Contacts	

Not local files on my device

What to do:

Periodic Audits of Access Controls:

- SharePoint
- OneDrive Teams
- Teams Channels
- Shared Calendars

Consider running a PowerShell script to purge old Sharing Links

Sensitivity Labels for Another Layer of Security

Protect targeted content with **Auto-Labeling**.

- Exchange
- SharePoint
- OneDrive

In the PIP example, create a custom SIT or Trainable Classifier to look for the Excel spreadsheet of a PIP and apply a label limiting access to specific security groups, like Human Resources.

Unsanctioned AI Example

A developer needs to review and debug confidential source code to help in product development but used an unsanctioned Generative AI.

What the developer inadvertently did was provide proprietary source code to a public generative AI available to anyone.

Competitors access this information, leading to loss of competitive advantage and potential legal battles over intellectual property.

proarch

Monitoring & Control AI Usage

Know What GenAI's Are in Use

Microsoft Defender for Cloud Apps can be configured to monitor generative AI user activity

It has 459 available cloud apps to choose with risk scores ranging from 10 (good) to 5 (not so good).

All of these are OFF by default.

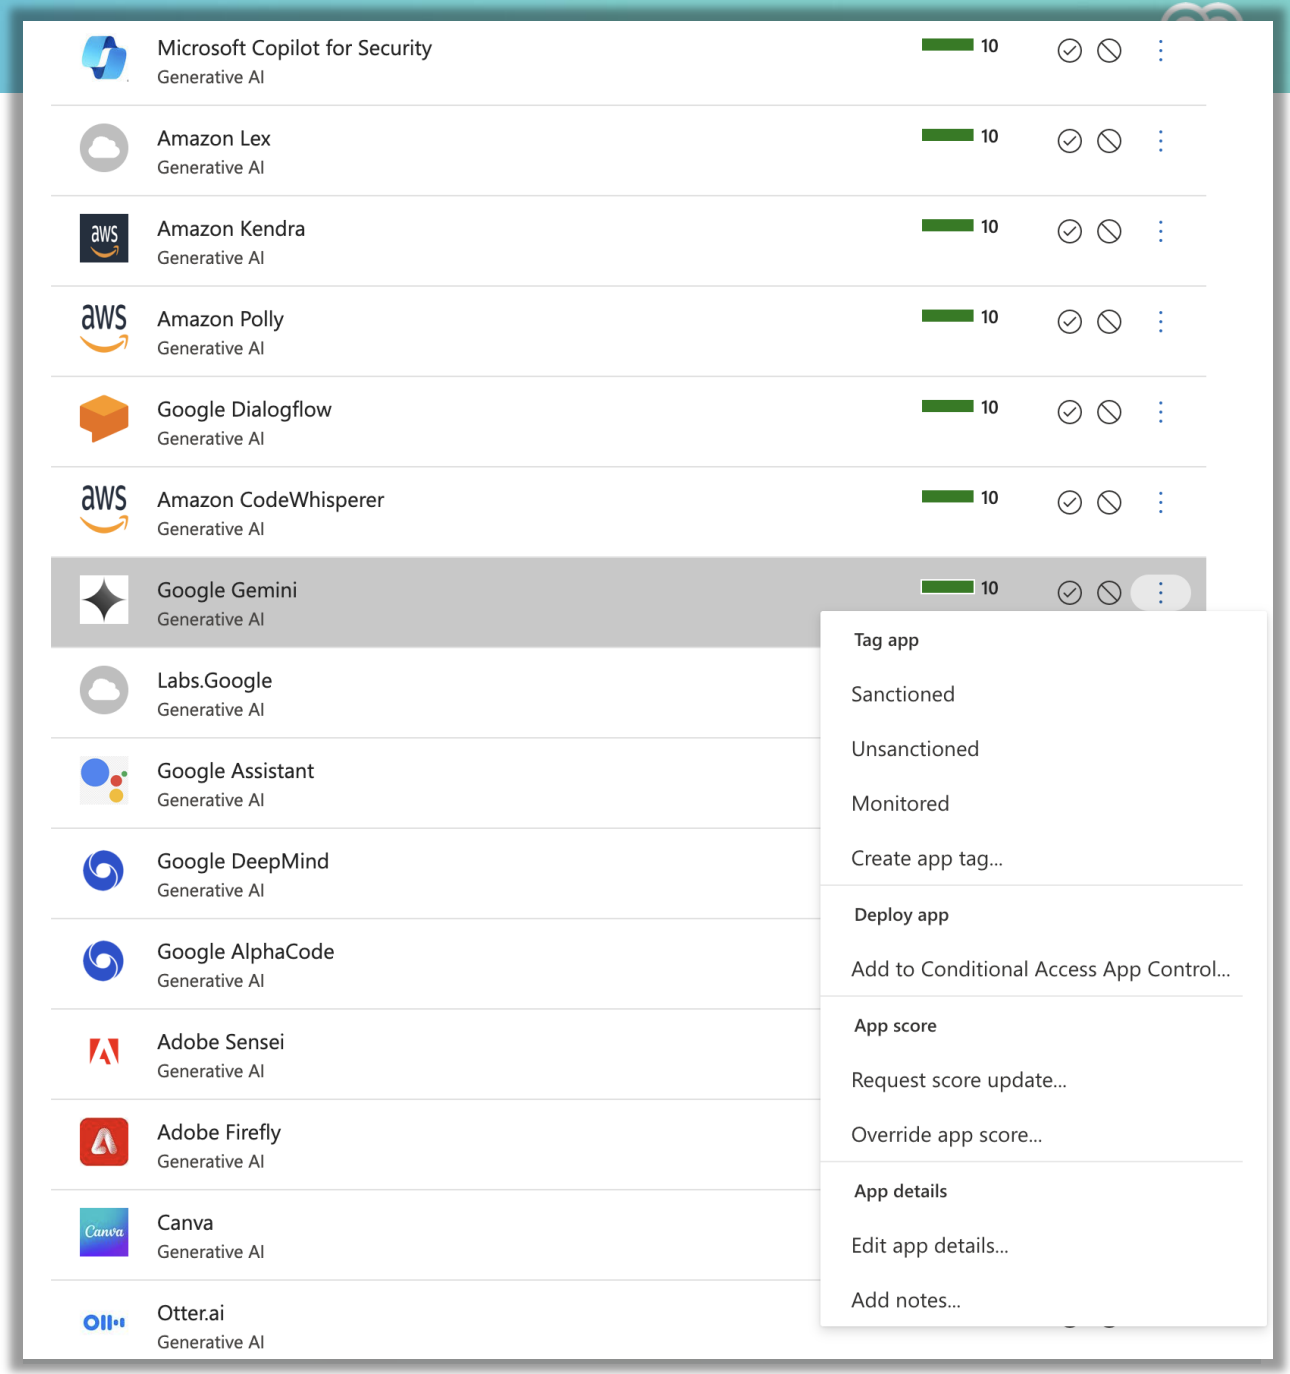

OPTION 1

Block all Unsanctioned Apps

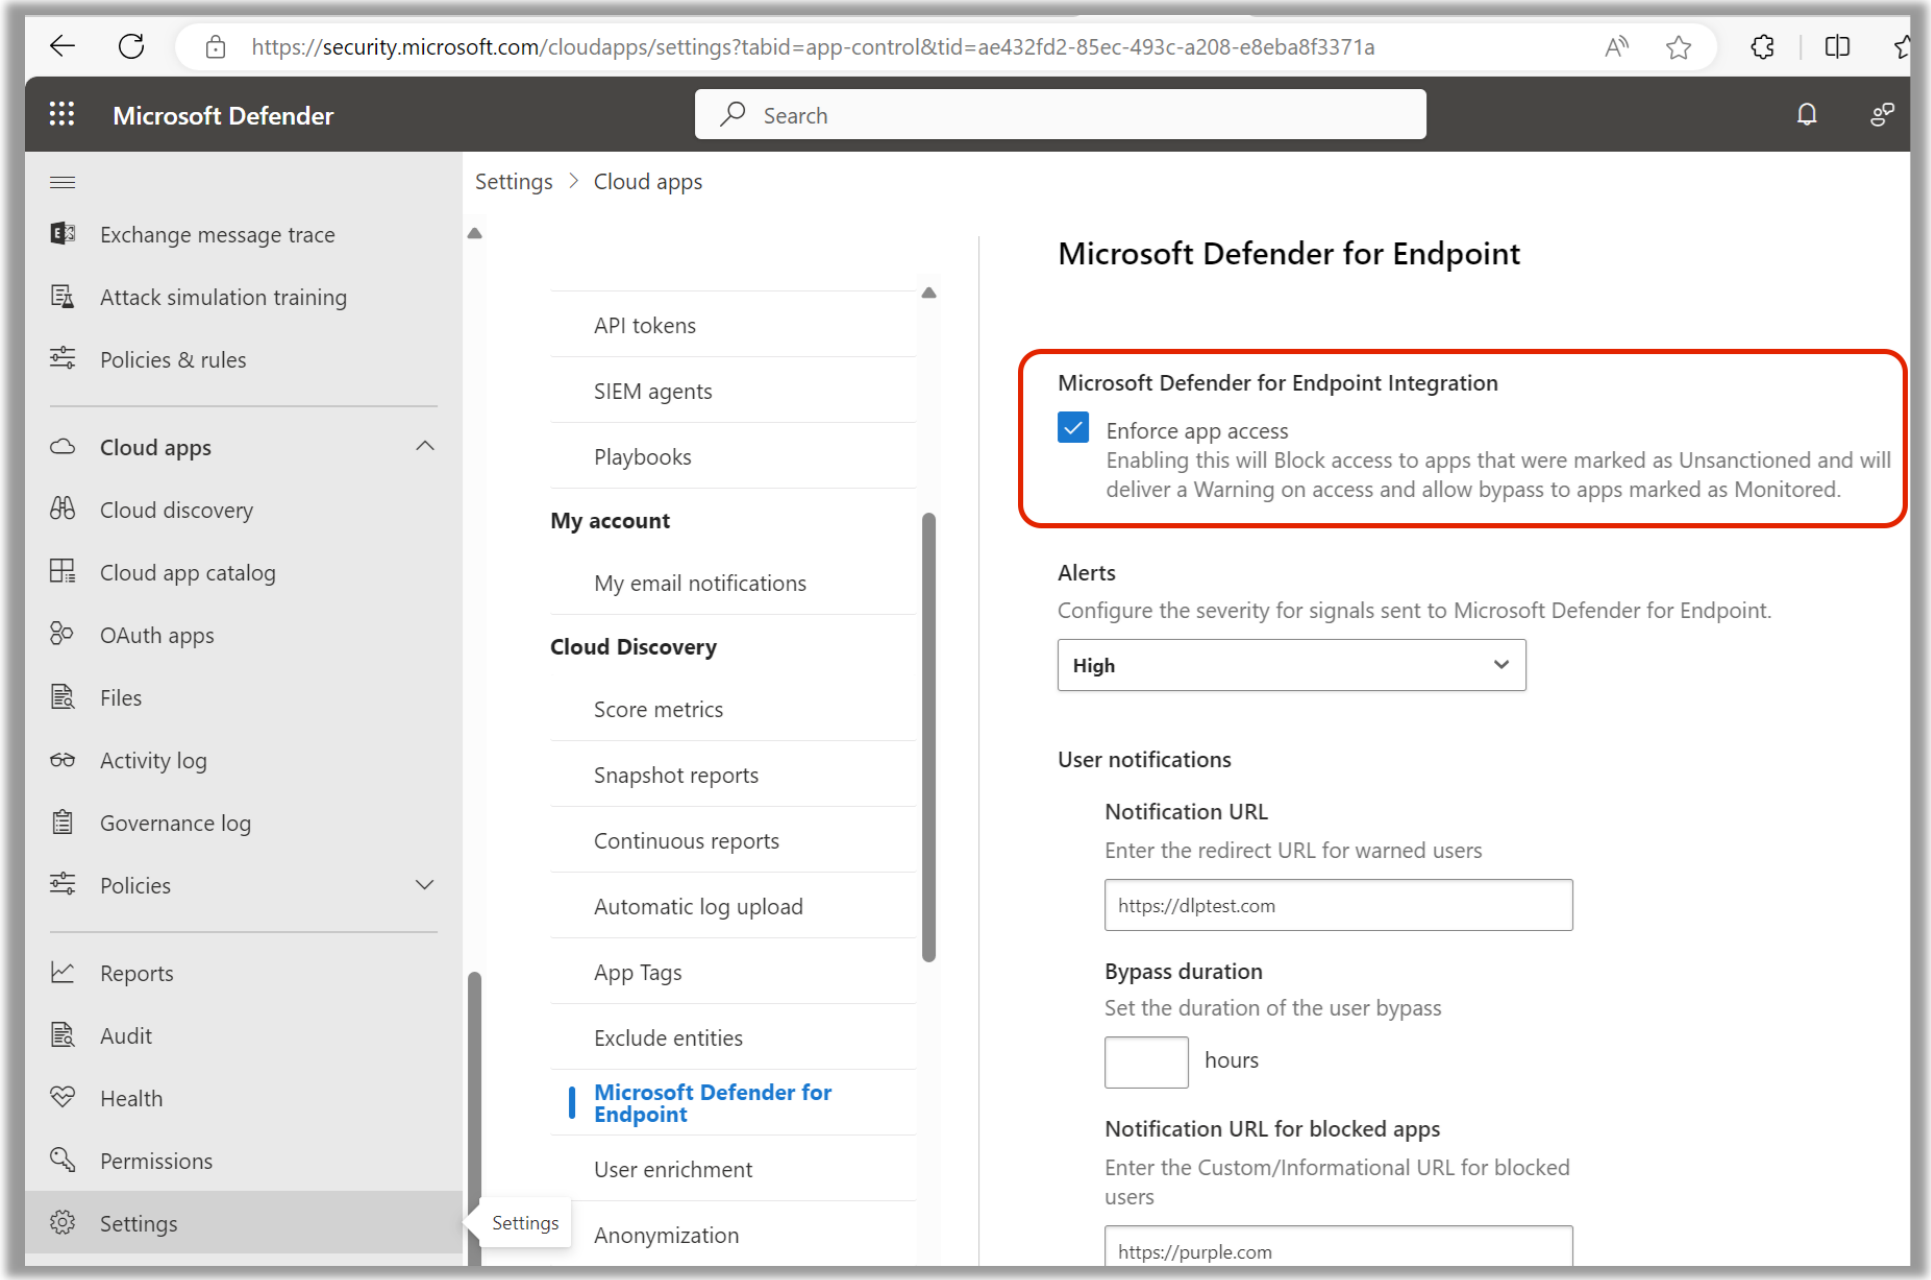

OPTION 2

Stopping Unsanctioned GenAI from Accessing Sensitive Data

1. Create a **Sensitivity Label** named “Internal Only”.
    - Configure the label to apply encryption and restrict access to only users within your organization
2. Apply the Sensitivity Label to the files either manually or label polices.
3. Create a new **DLP policy** and configure the condition for that Sensitivity Label and optionally other rules for similar content.
    - Within DLP Settings, configure a restricted Service Domain or Group to identify the URLs for unsanctioned generative AI (DLP is not integrated with Defender for Cloud Apps (formerly MCAS), so the URLs need to be re-entered).
    - Add the Action “Audit or restrict activities on devices” to the DLP rule with the Block action set.
4. Ensure that **domain restrictions** are applied to Microsoft Edge.
5. Configure unallowed browsers in DLP Settings.

AI Hub (preview)

Analytics and AI focused monitoring and protection policies

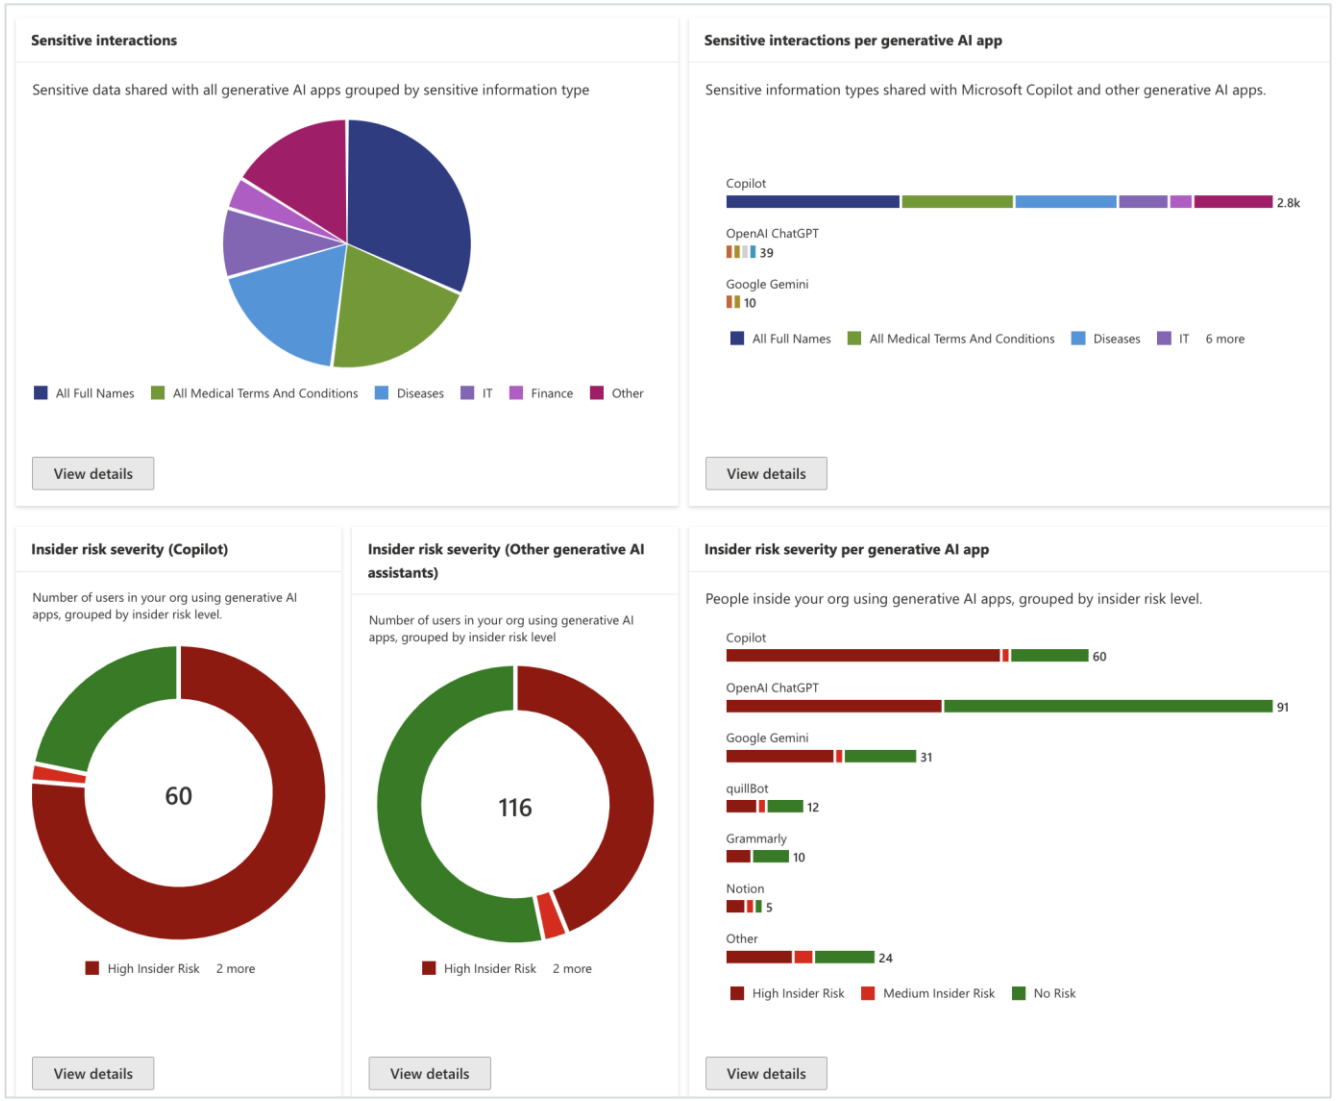

Automatic policy generation

Data security for AI

Data security risks can range from accidental oversharing of information outside your organization to data theft with malicious intent. Set up protection policies to manage your data security risks with AI assistants.

Here's what we'll set up for you:

Policy to be created

Adaptive protection to block sensitive information pasted or uploaded to 3P AI apps

Data loss prevention policy: **Microsoft AI hub - Adaptive Protection in AI assistants**

Uses Adaptive Protection to give a warn-with-override to elevated risk users attempting to paste or upload sensitive information to other AI assistants in Edge, Chrome, and Firefox. This policy covers all users and groups in your org in test mode.

Labels exist

Information protection policy for sensitivity labels

Information Protection

Sets up default sensitivity labels to preserve document access rights and protect Copilot output. Users can choose to label items in Outlook, Word, Excel, PowerPoint, and other locations.

proarch

Questions ?

THANK YOU FOR JOINING US | PROARCH.COM

Special offer! Free 2-hour data security session.
Contact your account team or letstalk@proarch.com.