Ready for Microsoft Copilot? Chances Are Your Data Isn't.

EBOOK

Copilot puts Microsoft Clippy on steroids

Microsoft Copilot is the next frontier many organizations will be exploring in 2025. An AI add-on to Microsoft 365, this new option promises to boost productivity and creativity, all while letting us further bow to our AI overlords of the future. But if Copilot is anything like Clippy, that paper clip thing everyone hated (except Darryl from "The Office"), then this is much to do about nothing, right? Most likely wrong.

Copilot looks to have some serious tricks up its sleeve. As a virtual assistant, Copilot will be able to quickly summarize content across any collection of data, like SharePoint sites, OneDrive, or Teams meetings. For Excel, it will be able to highlight trends and analysis across any data set with easy-to-leverage natural language queries.

It can integrate with CRMs and help auto generate customer proposals or details about accounts with ease. Even your PowerPoint presentations can be graphically fine-tuned for the audience you're delivering to.

Pretty cool, huh? Well, if you've got data protection on the brain, every situation listed above should be sending up red flags.

Oh, the secrets your data can tell

One of the main concerns around Copilot is the all–access pass it can bring to your tenant data. This would be fine if it were just a user and their own data, but permissions across your tenant suddenly become very important in the Copilot world.

Any data that is overpermissioned can quickly become fuel for the Copilot engine. If permissions aren't set right, acquisition plans, employee salaries, sensitive customer information, or medical records could be ingested and spit out by Copilot to unsuspecting and unprivileged users. Sure, users could probably stumble onto this data navigating OneDrive, but placing an AI–powered brain designed to search and destroy on top of your data creates a whole new level of crazy. Certainly, a recipe for disaster for IT teams looking to ensure good data protection hygiene across an organization's sensitive content.

What Microsoft Copilot Security can and can't do

So, how does Microsoft help you secure data with Copilot, and more importantly, where does it fall short?

The first thing to understand is Microsoft Purview Sensitivity Labels. With the ability for users to mark data as confidential or internal, these labels enable Microsoft to treat sensitive content with the respect it deserves. Copilot will also respect these labels and, in theory, block this data from being served up in prompts based on user permissions.

While sensitivity labels are a nice approach to data hygiene, they have their shortcomings. Specifically, the process is user-driven. If a user doesn't label data properly or doesn't even understand what sensitive data looks like, it's open for business in some regards for Copilot.

While Purview does have auto-labeling capabilities, it can cause problems with DLP accuracy when looking beyond Microsoft 365. More about that in the "Third-party data protection" section below.

Another important aspect to consider is data permissions, which again enables Copilot access to any and all data not permissioned right. This is always a struggle as users often prioritize oversharing in the spirit of collaboration, rather than implementing right-sized permissions based on data security needs.

For more Microsoft Copilot security info, you can **read this Microsoft Article**.

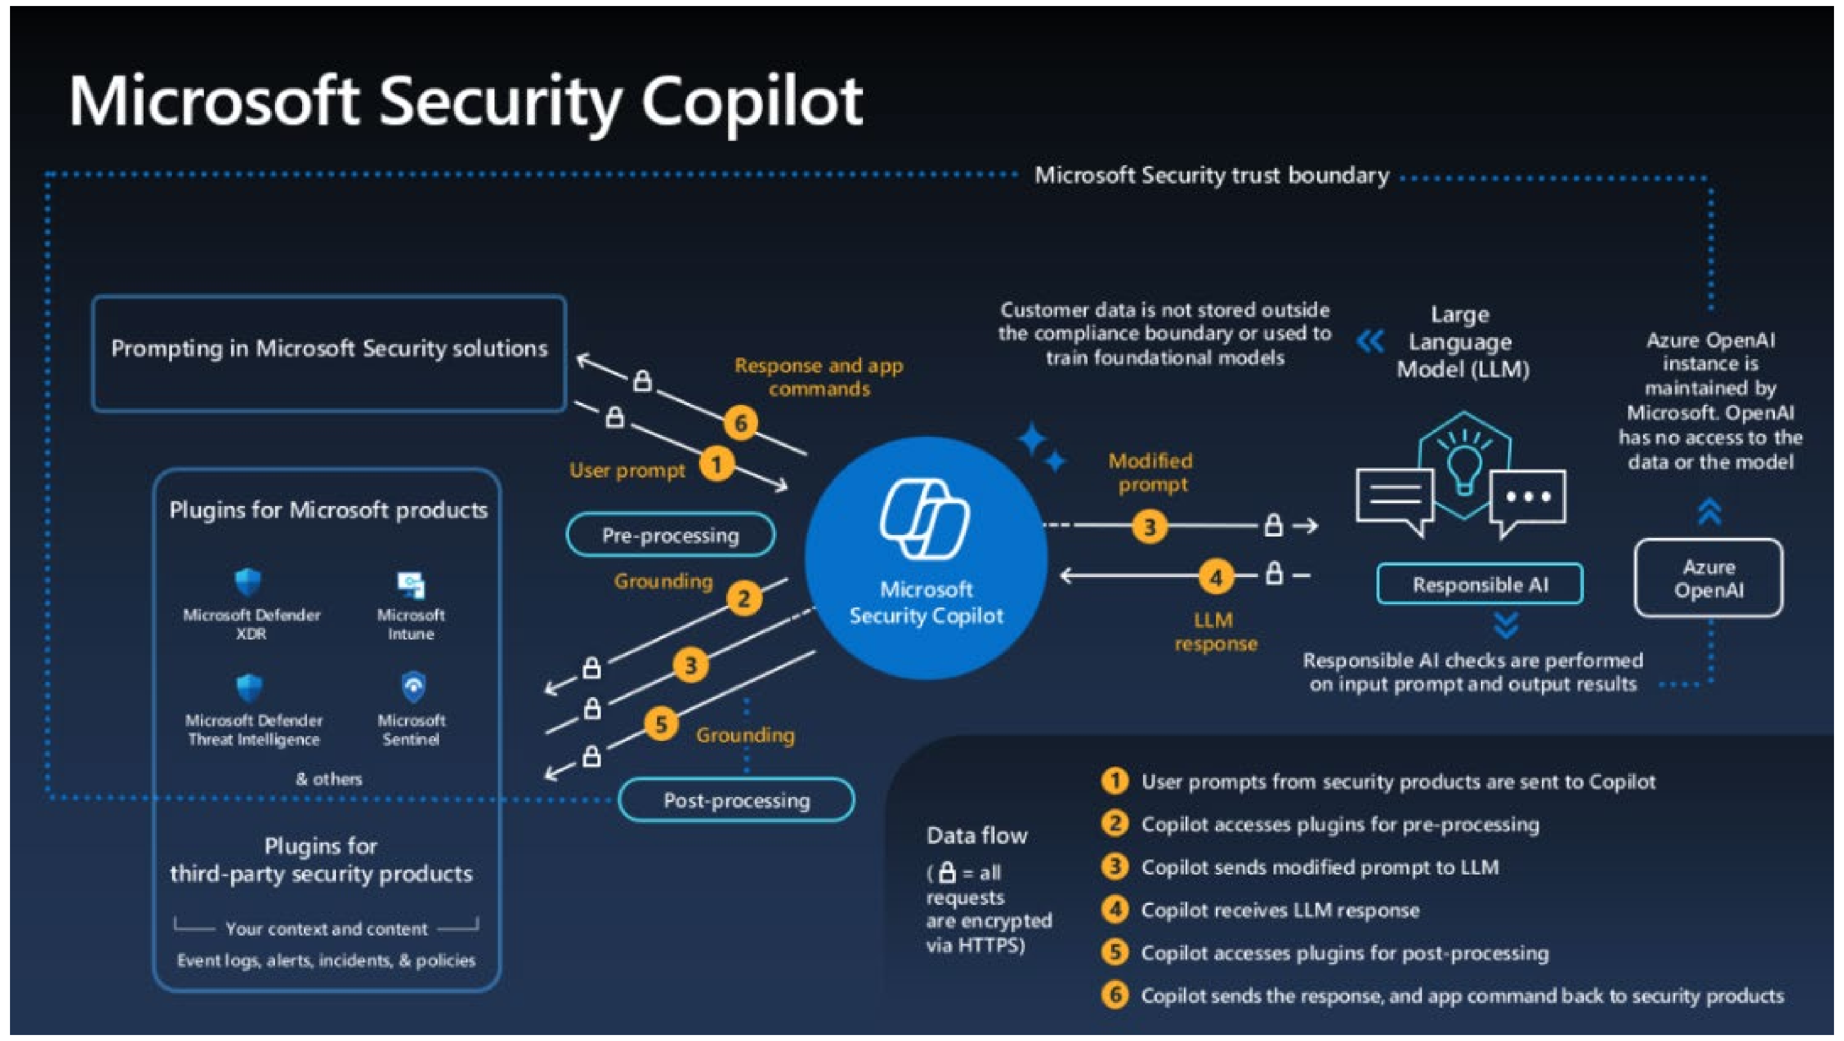
Microsoft Security Copilot
Microsoft Security trust boundary
Prompting in Microsoft Security solutions
Response and app commands
User prompt
Pre-processing
Grounding
Grounding
Post-processing
Plugins for Microsoft products
Microsoft Defender XDR
Microsoft Intune
Microsoft Defender Threat Intelligence
Microsoft Sentinel
& others
Plugins for third-party security products
Your context and content
Event logs, alerts, incidents, & policies
Microsoft Security Copilot
Modified prompt
LLM response
Customer data is not stored outside the compliance boundary or used to train foundational models
Large Language Model (LLM)
Responsible AI
Responsible AI checks are performed on input prompt and output results
Azure OpenAI instance is maintained by Microsoft. OpenAI has no access to the data or the model
Azure OpenAI
Data flow
(🔒 = all requests are encrypted via HTTPS)
1 User prompts from security products are sent to Copilot
2 Copilot accesses plugins for pre-processing
3 Copilot sends modified prompt to LLM
4 Copilot receives LLM response
5 Copilot accesses plugins for post-processing
6 Copilot sends the response, and app command back to security products

Microsoft Copilot security, and how third-party security integrates

Gartner Magic Quadrant for Security Service Edge

So, what new AI data protection tips and tricks should you be thinking about when evaluating and deploying Copilot? In no particular order, I give you the magical list that will make your Copilot journey much less painful:

- **Copilot prompt visibility**
  As you embrace Copilot, it's important to understand how your users are interacting with this new tool. The odd adage "you don't know what you don't know" rings true here. Visibility is king, and seeing input prompts to Copilot gives you important context as to how your users and data are getting along. In order to achieve individual user level prompt visibility, look for third–party solutions that support inline inspection, like that of **Gartner Magic Quadrant for Security Service Edge** (more on that later). This allows auditing and cataloging of all prompts sent to Copilot, which is the goal of this best practice.

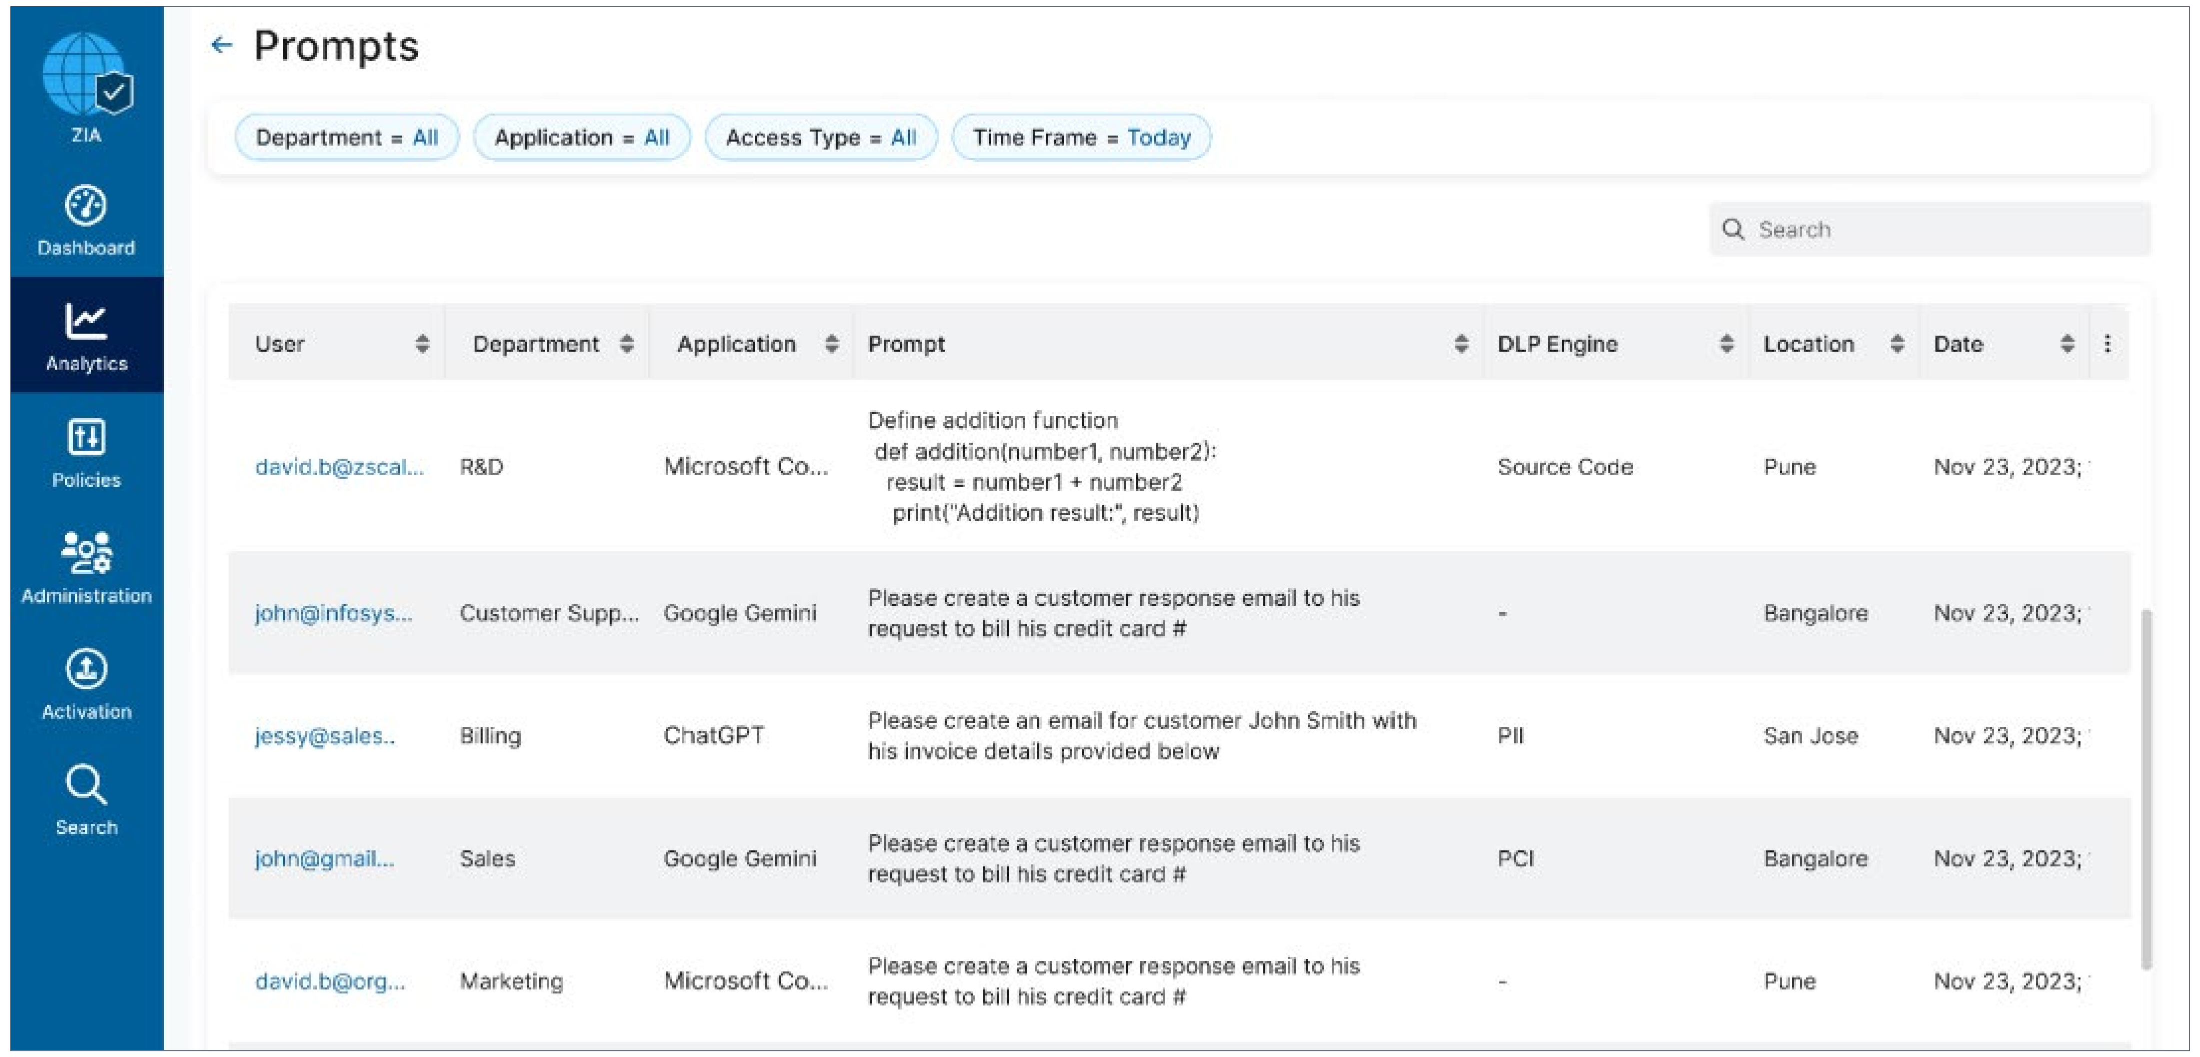

User	Department	Application	Prompt	DLP Engine	Location	Date

david.b@zscal...	R&D	Microsoft Co...	Define addition function
			def addition(number1, number2):
			result = number1 + number2
			print("Addition result:", result)
				Source Code	Pune	Nov 23, 2023;

john@infosys...	Customer Supp...	Google Gemini	Please create a customer response email to his request to bill his credit card #	-	Bangalore	Nov 23, 2023;
jessy@sales..	Billing	ChatGPT	Please create an email for customer John Smith with his invoice details provided below	PII	San Jose	Nov 23, 2023;
john@gmail...	Sales	Google Gemini	Please create a customer response email to his request to bill his credit card #	PCI	Bangalore	Nov 23, 2023;
david.b@org...	Marketing	Microsoft Co...	Please create a customer response email to his request to bill his credit card #	-	Pune	Nov 23, 2023;

User–level prompt visibility (from Zscaler Console)

- **Permission auditing via CASB and DLP**
  The next trick that is important is right out of the **cloud access security broker** (CASB) playbook. With a full fledged **data loss prevention** (DLP) engine, CASB allows you to troll through all your data in OneDrive and identify sensitive data that may have oversharing attached to it.

  Either with excessive internal sharing, or worse yet open internet links, CASB lets you audit sensitive files (as found by its DLP engine), and revoke excessive sharing. With CASB you can also set policy to control future risky sharing attached to sensitive data. All this helps you enforce proper permission levels across your sensitive data, and ensure Copilot has the correct rules of engagement before you deploy and use it.

- **Find and update missing Purview sensitive labels**
  As mentioned, Purview Sensitivity labels are a nice feature in the Microsoft ecosystem, but their effectiveness depends on correct usage. Users often fail to assign the right labels due to forgetfulness or difficulty identifying sensitive data.

  To address these issues, industry leading CASB's can find and update missing labels. By ingesting your organization's Purview labels and scanning OneDrive, CASB's can automatically update sensitive data with the appropriate missing sensitivity labels. The key to this process is again the DLP engine within the CASB, which uses extensive inspection via dictionaries to identify all types of sensitive data. In the Copilot world, this is a significant upgrade to data security, as these sensitivity labels can ensure right-sized output prompts based upon the user's permissions.

- **Close Copilot misconfigurations via SSPM**
  If you're not familiar with **SaaS security posture management** (SSPM), this an important step towards closing dangerous gaps. Built with a catalog of known misconfigurations, SSPM helps you scan SaaS Platforms, namely Copilot and Microsoft 365, and identify risky misconfigurations that may lead to data exposure.

  Scanning and remediating bad settings in your Microsoft deployment can be a life saver, especially over time. As your Copilot or Microsoft 365 deployment grows and adapts, the continuous assessment of SSPM ensures configuration drift doesn't re-introduce new risks to your data over time. From the screen shot below of Zscaler's SSPM, you can see the types of risks you can identify through this approach.

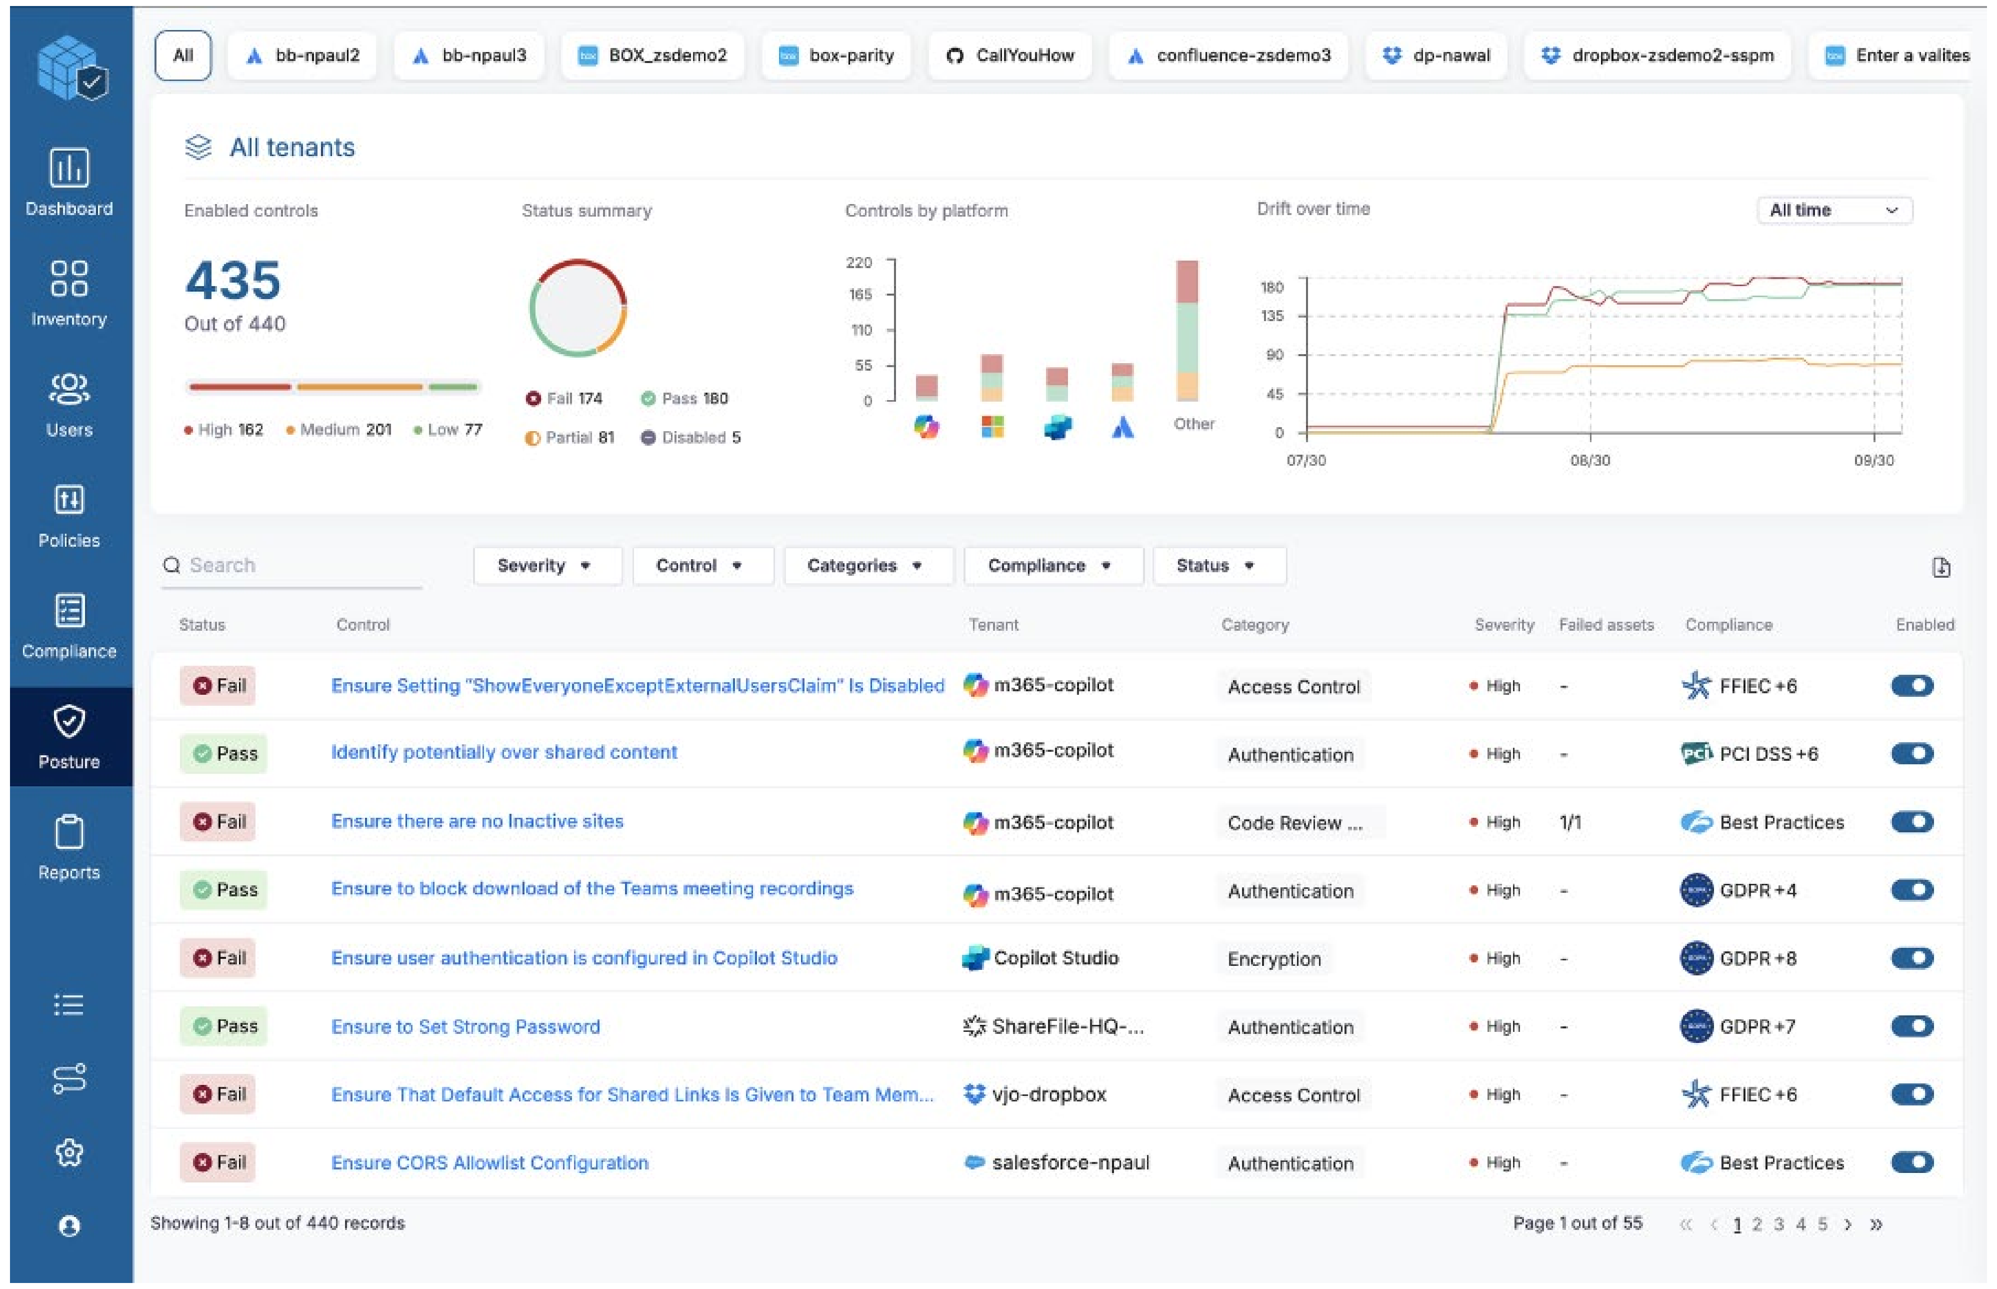

Status	Control	Tenant	Category	Severity	Failed assets	Compliance	Enabled
Fail	Ensure Setting "ShowEveryoneExceptExternalUsersClaim" Is Disabled	m365-copilot	Access Control	High	-	FFIEC +6	●
Pass	Identify potentially over shared content	m365-copilot	Authentication	High	-	PCI DSS +6	●
Fail	Ensure there are no Inactive sites	m365-copilot	Code Review ...	High	1/1	Best Practices	●
Pass	Ensure to block download of the Teams meeting recordings	m365-copilot	Authentication	High	-	GDPR +4	●
Fail	Ensure user authentication is configured in Copilot Studio	Copilot Studio	Encryption	High	-	GDPR +8	●
Pass	Ensure to Set Strong Password	ShareFile-HQ-...	Authentication	High	-	GDPR +7	●
Fail	Ensure That Default Access for Shared Links Is Given to Team Mem...	vjo-dropbox	Access Control	High	-	FFIEC +6	●
Fail	Ensure CORS Allowlist Configuration	salesforce-npaul	Authentication	High	-	Best Practices	●

Showing 1-8 out of 440 records

Page 1 out of 55 « ‹ 1 2 3 4 5 › »

Find dangerous Copilot misconfigurations (From Zscaler Console)

- **Block sensitive data headed to Copilot with inline DLP**
  Another important aspect of Copilot use is controlling what sensitive data is sent into Copilot. While Microsoft Copilot does not train on input data, allowing sensitive data into prompt workflows does increase risk.

  If users leverage internal information like credit cards or PII, this data can be reflected in the output, which users often accept as correct and valid. In this context, the sensitive data is now at much more risk to be leveraged in other ways as output by the user. To control this, having an inline DLP solution in place to monitor Copilot prompt transactions is best practice to ensure sensitive data doesn't become part of the Copilot workflow.

Why you need third-party data protection for Copilot

Now that we understand some best practices for securing Copilot data, how should you approach these steps, especially when these techniques are not native to the Microsoft ecosystem?

First, it's paramount to point out Data Loss Prevention (DLP) for its starring role in protecting sensitive data. DLP as a technology helps you inspect and classify sensitive data. Most Copilot best practices mentioned above revolve around a DLP engine. While you may be tempted to bring in individual DLP point products to tackle each Copilot issue separately, that's in reality a bad plan.

Point products, each with their own DLP engine, can lead to inconsistent alerts and increased complexity. The same data might trigger alerts in one engine but not another, making it hard for administrators to manage effectively. Customizing DLP policies would also require multiple configurations across different consoles, which can be a fast track to frustration.

Secondly, as you gain visibility into sensitive data, you'll likely want to make policy decisions about where it can go, who should have access, and what should not leave the organization. Your Copilot approach needs to be comprehensive and scalable, allowing you to add new protection channels as your data protection efforts evolve. Data can be lost through various channels, such as web uploads, emails, USB drives, SaaS sharing, IaaS misconfigurations, or other generative AI applications. Think of it as securing your digital fortress from all angles, not just the front door.

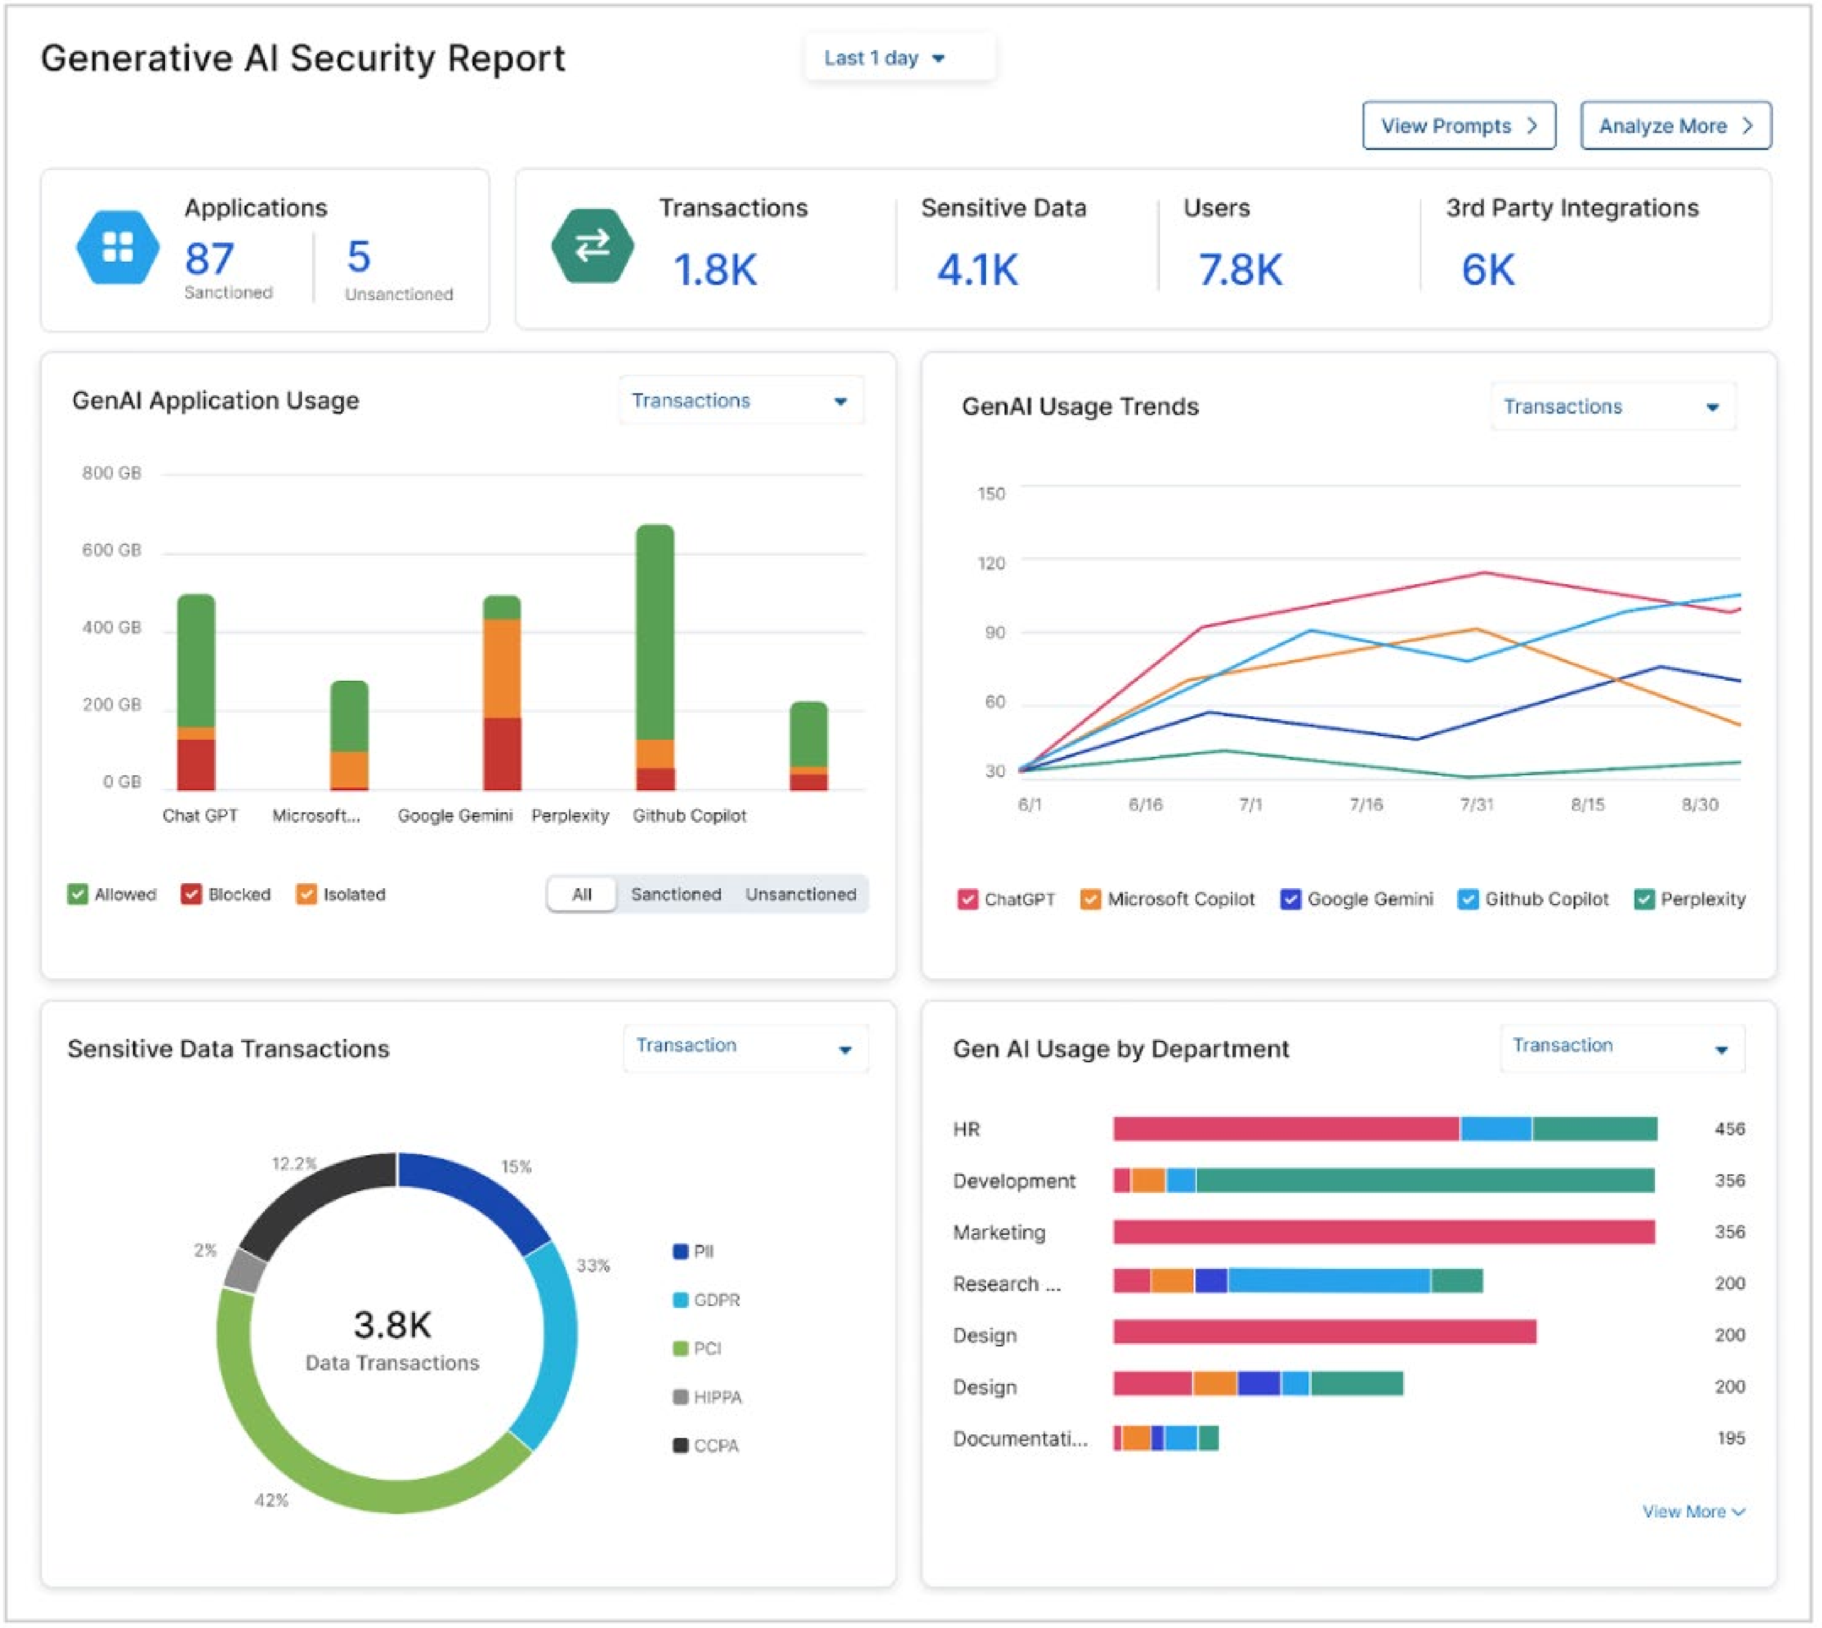

Generative AI Security Report

Last 1 day

Applications: 87 Sanctioned, 5 Unsanctioned

Transactions: 1.8K

Sensitive Data: 4.1K

Users: 7.8K

3rd Party Integrations: 6K

In-depth visibility of Microsoft Copilot and Public GenAI use (Zscaler Console)

Both these issues above lead us to the requirement of a Platform approach for data protection. This is where the **Gartner Magic Quadrant Security Service Edge** (SSE) comes in. Recognized and proven in the industry as the best architecture for a **data protection platform**, SSE helps organizations address all the concerns mentioned above. It ensures a flexible, cloud-delivered approach that can easily scale to meet future protection needs.

Gartner's SSE framework is designed to integrate all aspects of data protection into a unified architecture. It secures data in motion, data at rest in clouds, and across endpoints. This modern, cloud-based approach helps retire complex and costly legacy data protection methods, while enabling addition of new innovative approaches to data security. With full API and inline SSL inspection capabilities, Gartner's SSE architecture makes it easy to enforce robust security for Copilot, Microsoft 365, and other data loss channels.

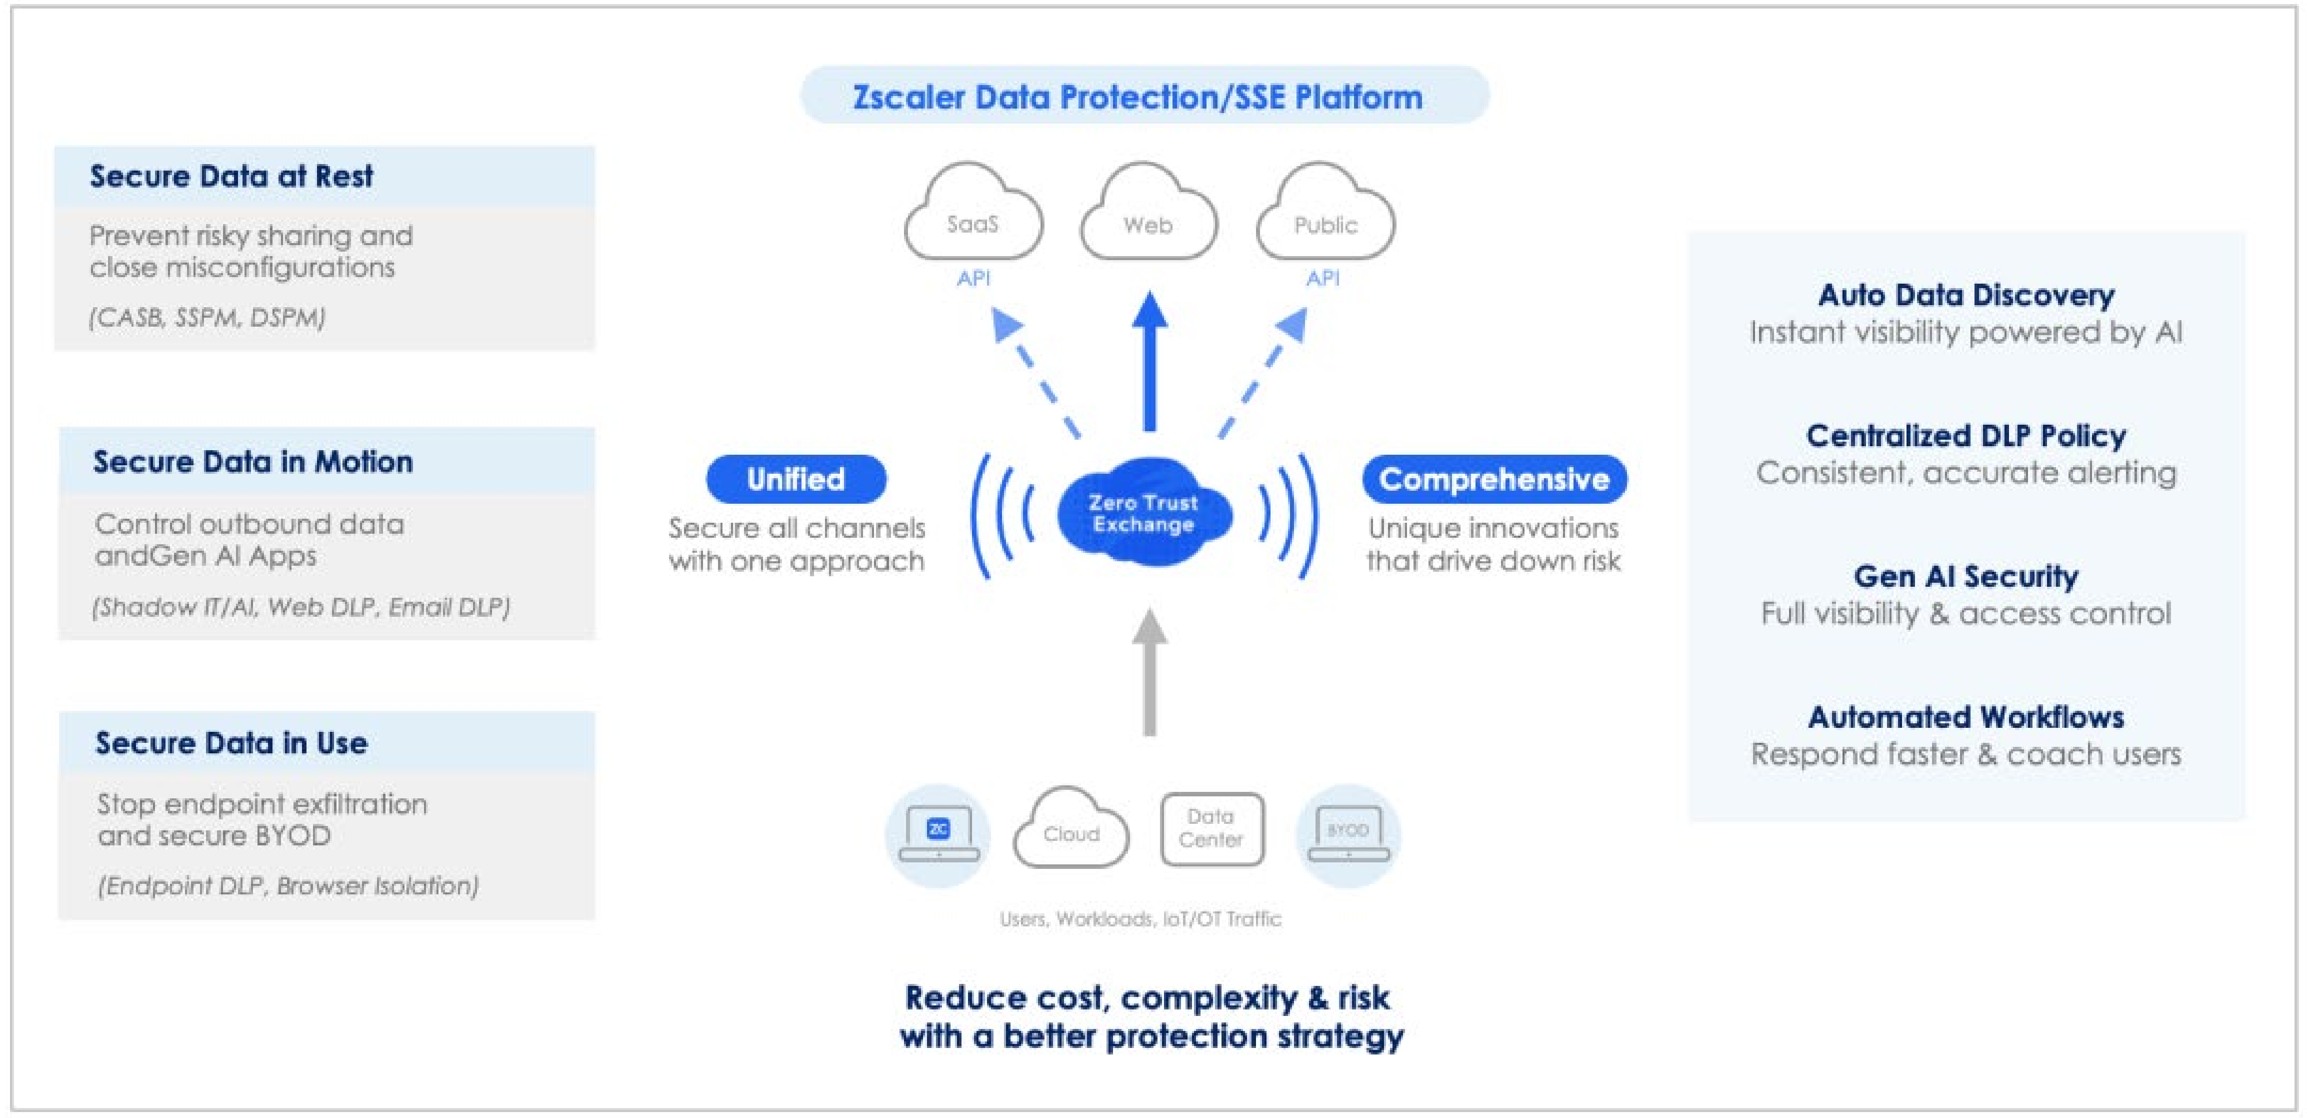

Zscaler Data Protection/SSE Platform

Secure Data at Rest

Prevent risky sharing and close misconfigurations

(CASB, SSPM, DSPM)

Secure Data in Motion

Control outbound data andGen AI Apps

(Shadow IT/AI, Web DLP, Email DLP)

Secure Data in Use

Stop endpoint exfiltration and secure BYOD

(Endpoint DLP, Browser Isolation)

Unified

Secure all channels with one approach

Comprehensive

Unique innovations that drive down risk

Auto Data Discovery
Instant visibility powered by AI

Centralized DLP Policy
Consistent, accurate alerting

Gen AI Security
Full visibility & access control

Automated Workflows
Respond faster & coach users

Reduce cost, complexity & risk with a better protection strategy

Zscaler's approach to data protection and SSE

Why Zscaler Data Protection for Copilot?

Why should you consider **Zscaler Data Protection** for securing Copilot? As a **leader in the Gartner Magic Quadrant Security Service Edge (SSE)** and the world's largest security cloud, Zscaler provides a unique and innovative approach to data protection. This has enabled **thousands of organizations like yours** to embrace a more secure future. With full integration support for Microsoft 365 and Purview sensitivity labels, Zscaler's Data Protection Platform helps organizations deploy a safer and more secure Copilot experience.

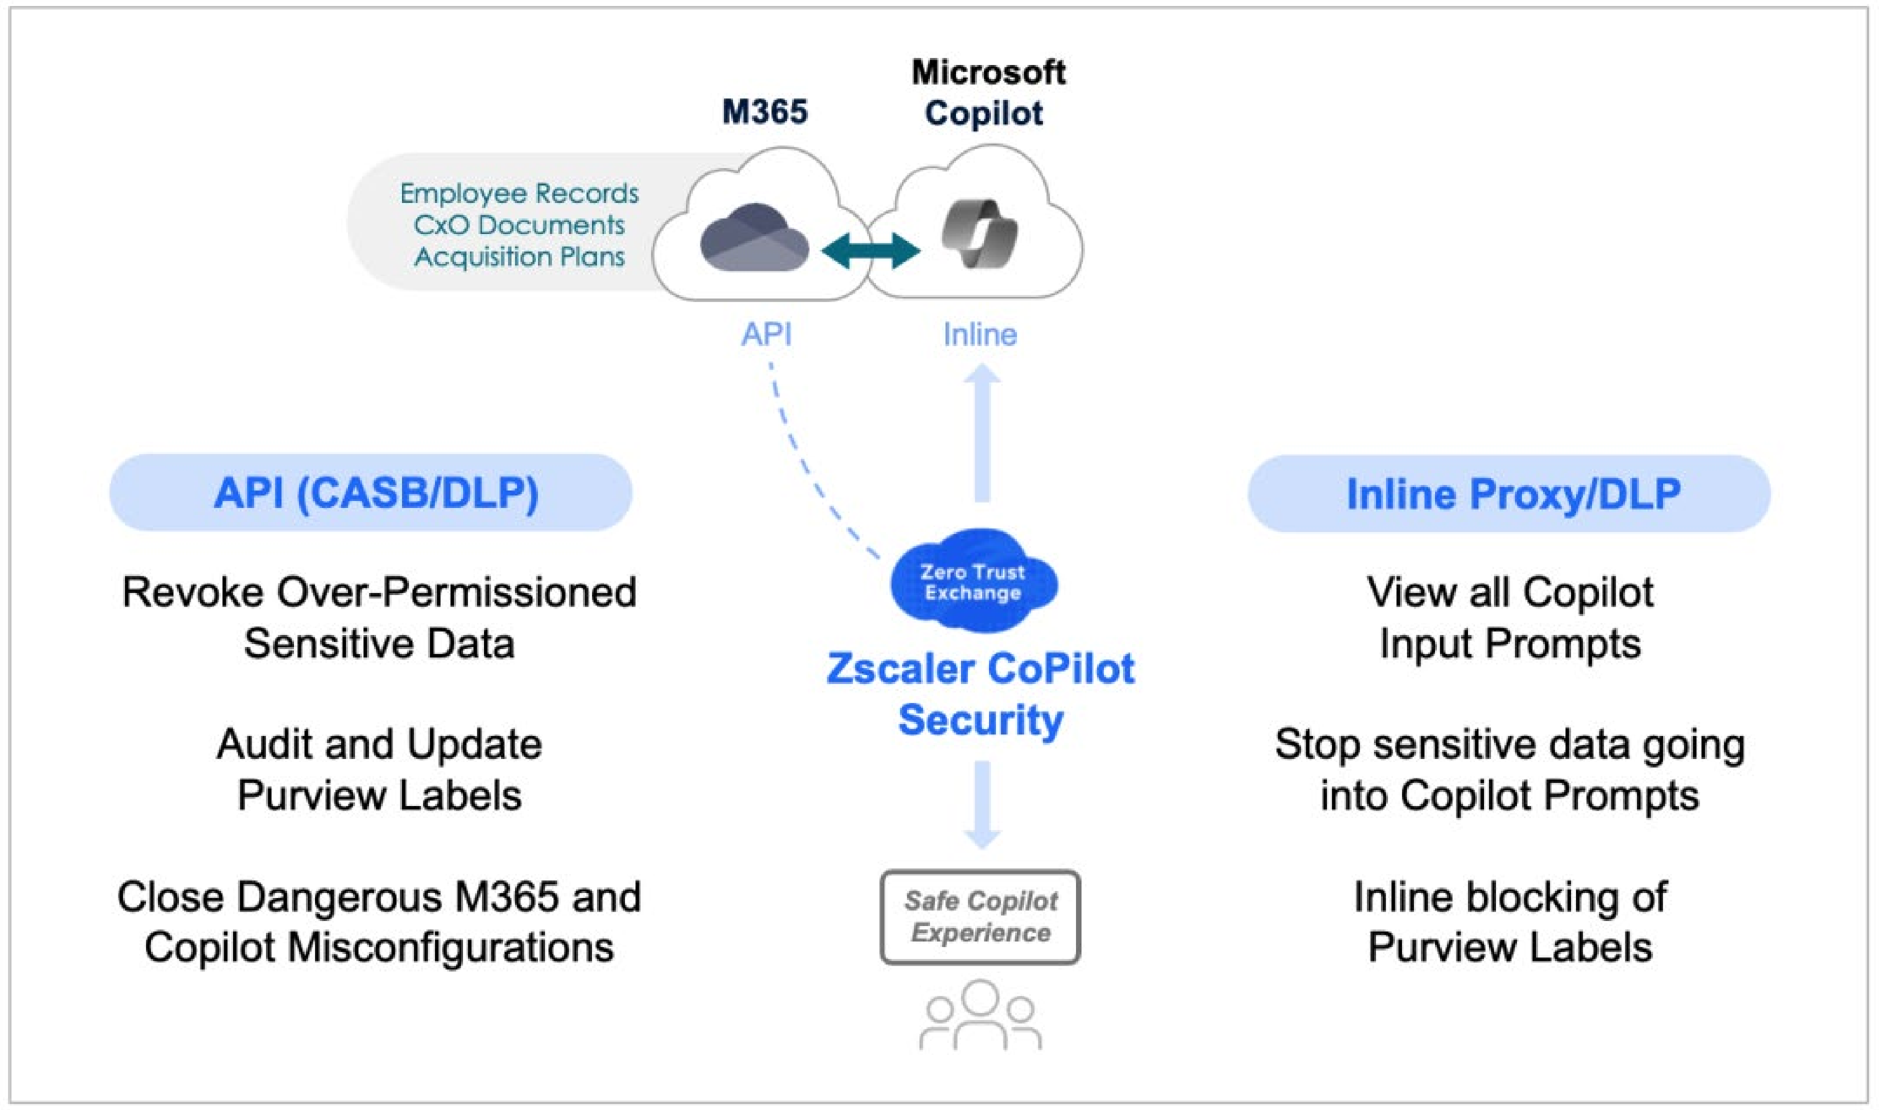
M365
Microsoft Copilot
Employee Records
CxO Documents
Acquisition Plans
API
Inline
API (CASB/DLP)
Revoke Over-Permissioned Sensitive Data
Audit and Update Purview Labels
Close Dangerous M365 and Copilot Misconfigurations
Zero Trust Exchange
Zscaler CoPilot Security
Safe Copilot Experience
Inline Proxy/DLP
View all Copilot Input Prompts
Stop sensitive data going into Copilot Prompts
Inline blocking of Purview Labels

How Zscaler secures Microsoft Copilot through API's and Inline

Here's how Zscaler Copilot Security can work with your Microsoft deployment for even better data security:

- **Revoke over-permissioned sensitive data**
  Use CASB/DLP to find and revoke access over-permissioned data that is open to Copilot access.
- **Audit and update Purview labels**
  Utilize APIs to update Purview tagging with missing labels to prevent Copilot access to sensitive data.
- **Close dangerous copilot misconfigurations**
  Use SSPM to identify and close risky Copilot or M365 configurations that expose data.
- **View Copilot user input prompts**
  See all user level inputs to Copilot for added context, visibility and incident response
- **Block sensitive data going into Copilot prompts**
  Via inline inspection, ensure that sensitive data is not sent into Copilot prompts or workflows.
- **Enforce inline blocking of Purview labels**
  Block sensitive data leaving the organization based on Purview labels using inline SSL inspection.

Remember, while your initial focus may be on Copilot data protection, Zscaler's platform helps you retire and consolidate legacy point products into a single, unified solution. You can scale your **CASB protection program** as you grow by easily adding new subscriptions like **inline control of web, email** and **GenAI apps**, **endpoint DLP**, **BYOD security**, and **public cloud security** (DSPM).

Learn more and next steps

If you're looking to deploy Copilot and need help ensuring Microsoft 365 data, watch the **Microsoft Copilot Readiness on-demand webinar** to learn more.

You can also **schedule a call** with us, or **request a demo**. We're here to help!

zscaler™ | **Experience your world, secured.™**