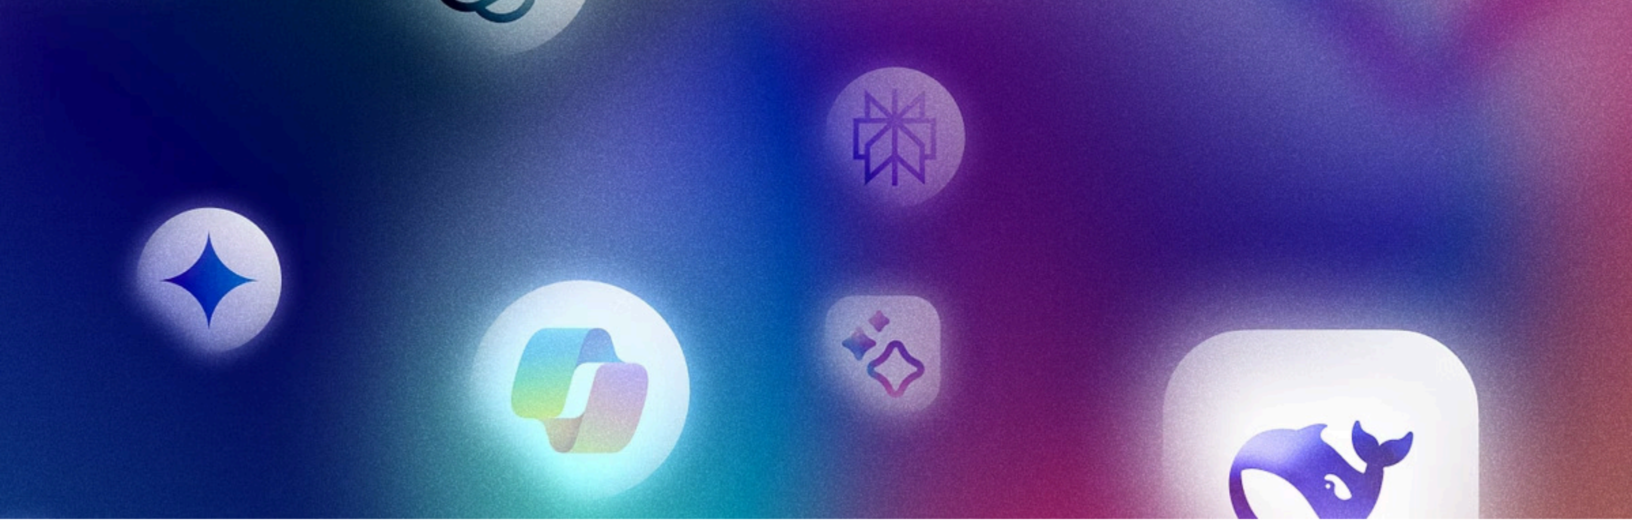

Market Guide

AI Discovery Methods Compared: Your Guide to Complete Visibility

The generative AI landscape is changing *fast.* In just the last year, Nudge Security has discovered **over 1,000 unique new AI tools hitting the market**, not to mention the MCPs and AI capabilities being added to virtually every other SaaS tool used by your workforce. Given this pervasiveness, it's critical to think about the full SaaS ecosystem when evaluating solutions to help you discover workforce AI use and mitigate risks.

But with a growing AI security market and a maze of vendor claims, how do you separate real AI visibility and control from smoke and mirrors? Let's cut through the noise and break down the most common AI discovery methods, so you can find the right fit for your organization.

This guide covers:

What is AI Discovery?

AI discovery is the process of identifying and cataloging all AI apps, accounts, integrations, user activities, and other dependencies that could expose your corporate data to third-party AI providers and beyond. An AI asset inventory includes AI apps that are procured and sanctioned by IT as well as "shadow AI"—the long tail of apps and embedded AI features that your employees experiment with and adopt *without* going through a formal approval process.

Not unlike the asset discovery and inventory systems developed for enterprise networks, the purpose of AI discovery is to create a comprehensive system of record of all AI assets introduced across your organization. But, instead of assets like servers, workstations, and installed software, your AI asset inventory captures third-party AI apps, user accounts, authentication and entitlement data, integrations across apps, MCPs, and AI supply chain dependencies.

Effective AI discovery can help you answer important questions like:

- Who is using AI across my organization?
- Have new (unapproved) AI tools been introduced outside of our governance process?
- Are AI tools integrated with apps that store our business critical data, and what level of access do these tools have?
- Are AI-enabled features in our other workplace technology solutions making our data accessible to LLMs?
- Is AI used in the supply chain of our other SaaS providers, and if so, what assurances have they provided on their data privacy practices?

Step one

Define the scope of what you need to discover. A common mistake is to focus only on discovering AI prompt activity in chatbots like ChatGPT, which is merely the tip of the iceberg given that AI is now making it's way into virtually every other SaaS tool used by your workforce.

Note: In the interest of avoiding scope creep, this guide focuses specifically on *workforce AI use*. In other words, the AI tools and embedded AI that can gain access to your sensitive data based on use by your employees. AI model security and agentic AI security are related topics, but warrant their own separate deep dive.

What Should AI Discovery Include?

When it comes to workforce AI use, here are the essential categories of AI assets you should consider for your discovery efforts:

Purpose-built AI apps

These are apps like ChatGPT, Perplexity, Claude, and others that offer a user interface for your workforce to interact with the LLM in order to answer questions, analyze data, get help writing, conduct research, generate images, and more. Ideally, you want to discover all user accounts for all apps in this category. Not an easy feat given that **the number of purpose-built AI tools is increasing exponentially**.

Tip: Look for solutions that can pattern match to recognize AI tools regardless of whether you've heard of them or not. Any static or vendor-managed list of AI tools will quickly become outdated.

Integrations that grant AI tools access to data

Many workforce AI tools offer easy connections to others platforms via OAuth grants or native marketplace apps. The ease of approving an OAuth grant or API key can entice users to hand over more access than they might realize—a well-meaning employee can inadvertently give AI tools access to, say, their entire corporate Google Drive, calendar, or email without realizing it. A recent example of this is **AI notetaker apps** that prompt users to grant access to calendars and contacts, while defaulting to "join every meeting." Suddenly, unapproved AI notetakers are popping up in meetings where sensitive topics are discussed, or where attendees have not consented to being recorded.

Tip: Look for solutions that can give you a robust inventory of app-to-app integrations through OAuth grants, APIs, and other methods, including details on scopes and potential risks, ideally with the ability to directly revoke risky access.

SaaS apps with embedded AI

The next category that is often overlooked, but is actually growing the fastest is SaaS apps that have embedded AI functionality within them. SaaS providers have been racing to launch AI-enabled features within their products and users have been equally eager to experiment with these new capabilities. However, enabling these embedded AI features may result in data access for the underlying LLM beyond what you want to allow.

Tip: Some SaaS and AI security solutions can surface and summarize the data training policies of your SaaS and AI vendors, including whether your data is used for model training, available opt-out options, retention periods, and other relevant information, saving you the time and effort of manually reviewing lengthy legal documentation.

MCP server integrations

We've seen a massive shift in how corporate data is accessed by AI tools following Anthropic's release of the open-source model context protocol (MCP). Major SaaS providers have followed suit, announcing their own MCP servers and native AI integrations. Instead of relying on one-off prompts and individual file uploads, these MCP servers radically expand data access by allowing LLMs to directly query SaaS apps through backend APIs.

While this expanded data set improves the quality and value of AI outputs by grounding them with your corporate data, it introduces significant new data security risks. MCP and other direct connections occur app to app, directly from the backend of providers like OpenAI to data stores like Salesforce via OAuth grants or APIs. And, they typically grant the AI tool the same data permissions as the user who created the connection, which can result in broader data access. So, in addition to the OAuth inventory mentioned above, discovering and governing API connections to AI tools is also critical in this new MCP-powered phase of AI use.

Tip: SaaS Security Posture Management (SSPM) tools can be helpful here as they typically inventory the integrations that have been enabled within the apps you have connected to the SSPM, so you don't have to manually review each app-to-app integration.

AI in the supply chain of your other SaaS providers

Another category that is often overlooked is AI in the supply chain of other SaaS app providers. An example of this would be companies who use AI tools to streamline customer support operations. Let's say you are a customer of Acme SaaS tools, and they use a third-party AI provider as part of their customer support process. If sensitive data from your environment is shared (say, tokens or API keys) as part of a support case, that data could now be accessible by the underlying LLM.

Tip: Ask your SaaS providers key questions about their third-party AI use, such as "Do you process our data through a third-party AI provider? If so, what security controls do you have in place for protecting our data privacy? What options do we have to limit what access could be accessible?"

AI user activities like chat prompts and file uploads

Last but not least, you'll want to consider what data or files users might be sharing with AI tools via chat prompts or file uploads. Typically, detecting these activities requires a browser-, endpoint- or network-based control that can recognize file-upload behavior or sensitive data being uploaded (like API keys or credit card numbers) and take action to alert the user of potential data security risks, notify the IT security team of this activity, or possibly intercept the prompt and obfuscate the sensitive data before it reaches the LLM.

Tip: Network-based controls are only useful to the extent that your employees work from within a "walled garden" visible to those controls. Browser-based controls, or a layered approach (which we'll cover below), tend to be more effective given the highly distributed nature of how and where modern work gets done.

Do You Need a New Tool for AI Discovery?

In most cases, this answer will be an emphatic "yes," but there are some aspects of AI discovery that you can get started with using tools you probably already own. Here are some of those methods:

Expense Reports and Financial Systems

If someone is paying for an AI tool, it will likely show up in your expense tracking systems at some point. While this data will show you paid tools, it will obviously miss free trials or free tier usage. Additionally, it might show you who the billing contact is, but you will still have a lot of digging to do to identify all the users of these paid tools.

Network Logs

Some teams set up rules within their network monitoring tools to look for signs of AI usage based on a specific list of known AI domains, or heuristics like domains ending in ".ai." This method provides visibility into the actual users of an AI tool and usage frequency, but the downside is you will have to either know what AI tools to look for, or be ready for some manual effort to separate the actual AI tools from the false positives.

IdP OAuth Grants

If your employees commonly use the "sign up with Google" or "sign up with Microsoft" options when they create accounts for new apps, you'll be able to see these by reviewing the OAuth grants that have been established via your IdP. This is helpful, but of course only works if the user signed up with their work email, and opted for this option rather than a username and password.

SSO Activity

Obviously, your SSO platform can only show you accounts and activity for the AI apps that have actually been onboarded into SSO, but this can be useful if you are trying to understand which AI tools are used most often and by which users and track the progress of AI adoption efforts. But, when it comes to uncovering unsanctioned AI use, this method will not be helpful.

Individual App Integrations

For your most critical apps (i.e. those in scope for your compliance efforts), you could go into the integration settings for each app to review OAuth grants and API connections that have been enabled for signs of connections with AI tools. Again, this would involve manual effort and regular audits, which limits the scalability and timeliness of the data.

The DIY Summary

The obvious benefit of using the AI discovery methods outlined above is avoiding the expense of adding another tool to your tech stack, but the data you can gather is limited, likely to become quickly outdated, and will require significant manual effort. Unless you operate in an environment with very strict controls over incoming and outgoing connections, you should likely consider a more complete and scalable solution.

Tools to Discover Workforce AI Use

Given the challenges outlined above, it's no surprise that the marketplace of security solutions is evolving rapidly to address the emerging risks posed by workforce AI use. Given the myriad of vendors and claims, it's important to understand the types of tools that are emerging to help with this challenge and the nuances of what they can and can't do.

At a high level, there are essentially two categories of tools to consider:

- **SaaS security solutions:** these platforms can discover use of shadow SaaS, and have expanded their capabilities to discover and categorize AI use, along with providing other AI security and governance capabilities.
- **AI security point solutions:** a new and rapidly expanding list of tools designed specifically for governing workforce AI use and protecting sensitive data.

Let's start with **SaaS security solutions** as there are differences in *how* these tools actually discover AI use that are important to understand. SaaS security solutions rely on one (or more) of the following methods to discover SaaS and AI use:

Network-based monitoring and/or endpoint agents

Browser-based discovery

Email-based discovery

API connections with SaaS apps

AI Discovery with SaaS Security Tools

Method 1: Network-based monitoring and endpoint agents (for example, CASBs)

These approaches rely on capturing user traffic—either via desktop agents installed on individual devices or network-level monitoring through firewalls, VPNs, or Cloud Access Security Brokers (CASBs). Similar to the network log monitoring option in the "DIY" section, this method detects and logs access to domains, providing visibility into what apps employees are using based on their web traffic.

Network-based monitoring and CASB solutions can theoretically provide broad visibility into AI tool usage, but in practice they face significant challenges in modern work environments. These methods struggle with decentralized workforces, BYOD policies, and the technical limitations of detecting modern AI applications with dynamic domains and encrypted traffic.

How it works:

Traffic inspection via network-level monitoring or desktop agents that flags traffic to the domains of known AI providers. Forward-looking only.

Pros:

- Monitors cloud service access
- Applies security policies

Cons:

- Limited to corporate networks
- Challenging for remote work
- Resource-intensive
- Requires endpoint agents
- No historical discovery

Questions to ask of vendors:

- Do you offer built-in rules to identify AI use? If so, how frequently are those rules updated?
- What is the average deployment time for your solution? What is the average time to value?
- Does your solution offer both inline and out-of-band (API-based) discovery and security monitoring?
- Does your solution detect app-to-app integrations, including MCP connections?

Method 2: Browser-based AI discovery

Browser-based AI discovery monitors app usage directly through a lightweight browser extension deployed to corporate devices. Extensions installed in corporate browsers can detect AI web visits, account signups, login activity, authentication methods, password strength, app usage patterns, file sharing, prompt content, and other rich behavioral and risk insights.

How it works:

Monitors for browser activity happening between users and the domains of AI tools.

Pros:

- Granular user activity data
- Real-time interventions
- Usage pattern insights

Cons:

- Misses mobile/personal device usage
- Limited historical data
- Depends on installation
- May have limited domain support

Questions to ask of vendors:

- Can you differentiate domain visits from login data?
- What metadata does the browser extension collect?
- How are you identifying activity related to AI tools vs. other SaaS apps?

Method 3: Email-based AI discovery

This approach analyzes corporate email communications from SaaS and AI providers for evidence of AI activities—things like welcome emails, password resets, billing notifications, MFA prompts, and security alerts. This discovery method works via read-only API connections to your IdP provider (like Google Workspace or Microsoft 365) in order to scan for emails related to AI app usage, identifying both the sanctioned and unsanctioned tools users have signed up for.

How it works:

Connects to your email provider and analyzes historical and future email activity for signs of AI use and captures rich context related to this activity.

Pros:

- Broad discovery coverage
- Finds unknown applications
- Historical insight
- Detects various account types, including username/password

Cons:

- Limited visibility into personal accounts
- Actual capabilities vary by vendor

Questions to ask of vendors:

- What metadata does your method of email discovery collect about app usage? (beware of solutions that are limited to email headers)
- Can you discover other details about AI use, like spend?

Method 4: API connections with SaaS apps (for example, SSPM solutions)

SaaS Security Posture Management, or SSPM solutions typically work via direct API connections with specific SaaS apps and can provide detailed visibility into app-to-app integrations, and app configurations. For example, if you connect Salesforce to your SSPM provider, you could then potentially see all of the app-to-app connections (via OAuth or API) that have been granted between Salesforce and any other SaaS tool, including AI tools.

However, these tools do not offer true discovery given that you have to establish integrations between the SSPM tool and the apps you want to monitor with it. While this method can't discover "shadow" AI use, it is helpful for gaining visibility into integrations between your critical SaaS apps and AI tools so you can audit and revoke data sharing entitlements with AI tools.

How it works:

API connection between the SSPM solution and supported SaaS apps within your environment.

Pros:

- Visibility into API and OAuth connections between AI tools and critical apps
- Enables review of data sharing entitlements for AI tools
- Continuous security monitoring for misconfigurations and risks

Cons:

- Limited to known apps
- API availability varies by SaaS provider
- Higher effort to deploy relative to other options
- Finite app coverage

Questions to ask of vendors:

- How does your solution handle remediation of security findings when API automation is not available?

Summary: SaaS Security Solutions

As you've likely surmised, no single approach for AI discovery can give you the full picture of AI use, which is why most SaaS security solution providers offer a layered approach comprising more than one discovery method.

For example, **Nudge Security** combines email discovery, a browser extension, SSO integrations, and API connections in order to address most AI discovery needs.

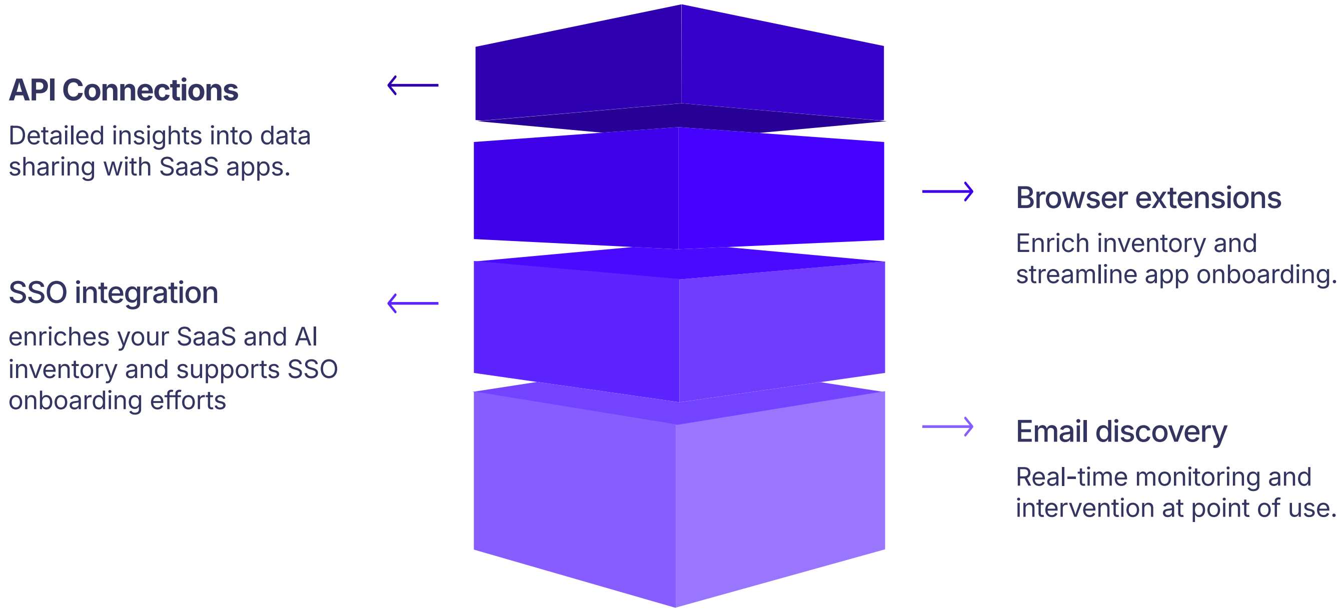

API Connections

Detailed insights into data sharing with SaaS apps.

Browser extensions

Enrich inventory and streamline app onboarding.

SSO integration

enriches your SaaS and AI inventory and supports SSO onboarding efforts

Email discovery

Real-time monitoring and intervention at point of use.

Emerging AI Security Tools

Now that we've covered how SaaS security solutions are evolving to meet the AI discovery and governance challenge, we'll circle back to the growing list of AI security point solutions.

For the most part, these tools leverage discovery methods similar to the browser-based discovery method described above. Some offer additional governance options like obfuscating sensitive info that users share in prompts (credit card numbers, SSNs, API keys, etc.)

The most important caveat to be aware of with these tools is that they are focused specifically on AI tools, not broader SaaS usage. And, as we've covered above, the lines are becoming very much blurred between what is an "AI tool" and what is a "SaaS tool" as SaaS vendors race to incorporate AI-enabled functionality. Additionally, integrations between SaaS tools and AI tools, as well as MCP servers introduce data security risks far beyond what a use might share in a prompt.

How AI security point solutions work:

AI security point solutions deploy advanced monitoring systems through browser extensions, endpoint agents, and API integrations with leading AI platforms to track and control how employees interact with AI tools.

Pros:

- Deep visibility into prompt-level interactions
- Real-time monitoring of specific AI tools
- Specialized AI risk detection capabilities
- Content filtering and sensitive data masking
- AI-specific policy enforcement options
- Protection against prompt injection attacks

Cons:

- Limited support of AI tools
- Fragmented visibility across the AI ecosystem
- Often miss AI embedded within SaaS apps
- Minimal historical data collection capabilities
- Implementation complexity for distributed workforces
- Separate from broader SaaS security governance
- Focus on symptoms rather than comprehensive governance
- Difficulty scaling with rapid AI tool proliferation

Questions to ask of AI-specific point solution vendors:

1. What specific AI security risks does your solution address? (prompt injection, data leakage, etc.)
2. How does your solution discover AI tools across the organization?
3. Can your solution detect AI embedded within existing SaaS applications?
4. How do you handle historical AI usage data versus forward-looking discovery?
5. What visibility do you provide into AI-to-SaaS integrations and data flows?
6. How does your solution balance security with employee productivity?
7. What governance capabilities do you offer beyond technical controls?
8. How does your solution integrate with our existing security stack?

Critical Questions to Ask of All Vendors

We've covered a lot of points to consider. To help you with your evaluation process, here's a condensed list of questions to ask of vendors regarding their AI discovery capabilities to assess how well they will meet your needs.

1

Discovery Methods

1. What AI discovery method(s) do you use?
2. Is your AI discovery method forward-looking only or can you discover AI assets created before your solution was deployed?
3. Does your AI discovery method rely on a known list of AI tools or can it dynamically identify new, never-before seen AI tools?

2

Asset Detection Capabilities

Which assets can your discovery method detect?

- AI apps
- User accounts
- OAuth integrations between AI tools and other apps
- API integrations between AI tools and other apps
- AI in the supply chain of other SaaS providers
- Data shared via AI prompts
- Trends of AI tool adoption and usage patterns
- Spend on AI tools

3

Insights and Intelligence

What insights can you provide on the AI tools that are discovered?

- Security program details
- Data training policy summaries
- Breach histories
- Data locality
- Compliance attestations

4

Integration Discovery

1. Can you discover integrations between AI tools and your other business tools? If so, how?
2. Can you discover MCP server integrations?

AI Discovery Methods Comparison Chart

As you've probably concluded, there's no magic bullet when it comes to finding all the AI tools and integrations floating around your company. Each discovery approach has its strengths and blind spots. Understanding what works best will help you pick the right solution (or mix of solutions) to keep tabs on your AI landscape.

Method	Pros	Cons
DIY AI Hunting	• Free (well, sort of)	• Time-consuming • Quickly outdated • Incomplete answers spread across multiple sources
Network-based monitoring and endpoint agents (e.g. CASBs)	• Monitors cloud service access • Applies security policies	• Limited to corporate networks • Challenging for remote work • Resource intensive • Requires network proxies or endpoint agents • No historical discovery
Browser-based discovery*	• Granular user activity data • Real-time interventions • Insights into usage frequency	• Misses mobile and personal devices • Historical visibility is limited to browser history • Requires browser plug-in • Discovery may be limited to "known" AI domains
Email-based discovery	• Broad discovery coverage • Finds unknown applications • Historical insight	• Limited visibility into personal accounts • Actual capabilities vary by vendor
API Connections	• Visibility into integrations between AI tools and critical apps • Enables review of data sharing entitlements for AI tools • Continuous security monitoring for misconfigurations and risks	• Limited to known apps • API availability varies by SaaS provider • Higher effort to deploy relative to other options • Finite app coverage

*This is the method typically employed by AI-specific point solutions

The Ideal AI Discovery Solution

Effective AI discovery is crucial for modern organizations managing a distributed workforce and complex environments. Because each discovery method has its strengths and limitations, organizations should carefully evaluate their specific needs, infrastructure, and work environment when selecting an AI discovery solution.

The ideal solution should not only discover AI applications but also uncover where AI is embedded and integrated across the entire SaaS ecosystem. As the AI landscape continues to evolve and SaaS providers add AI-enabled capabilities, the line between what constitutes and "AI tool" vs. a "SaaS tool" will continue to dissolve.

By understanding the various discovery methods available and their respective trade-offs, teams can make informed decisions about which approach—or combination of approaches—best suits their needs now and can scale for the future. Remember that AI discovery is not a one-time exercise but an ongoing process that requires continuous monitoring and adaptation to keep pace with the evolving nature of AI use.

AI Discovery with Nudge Security

Nudge Security users a layered approach that combines multiple discovery methods, with email-based discovery at its core, to provide unmatched visibility into AI (and SaaS) use.

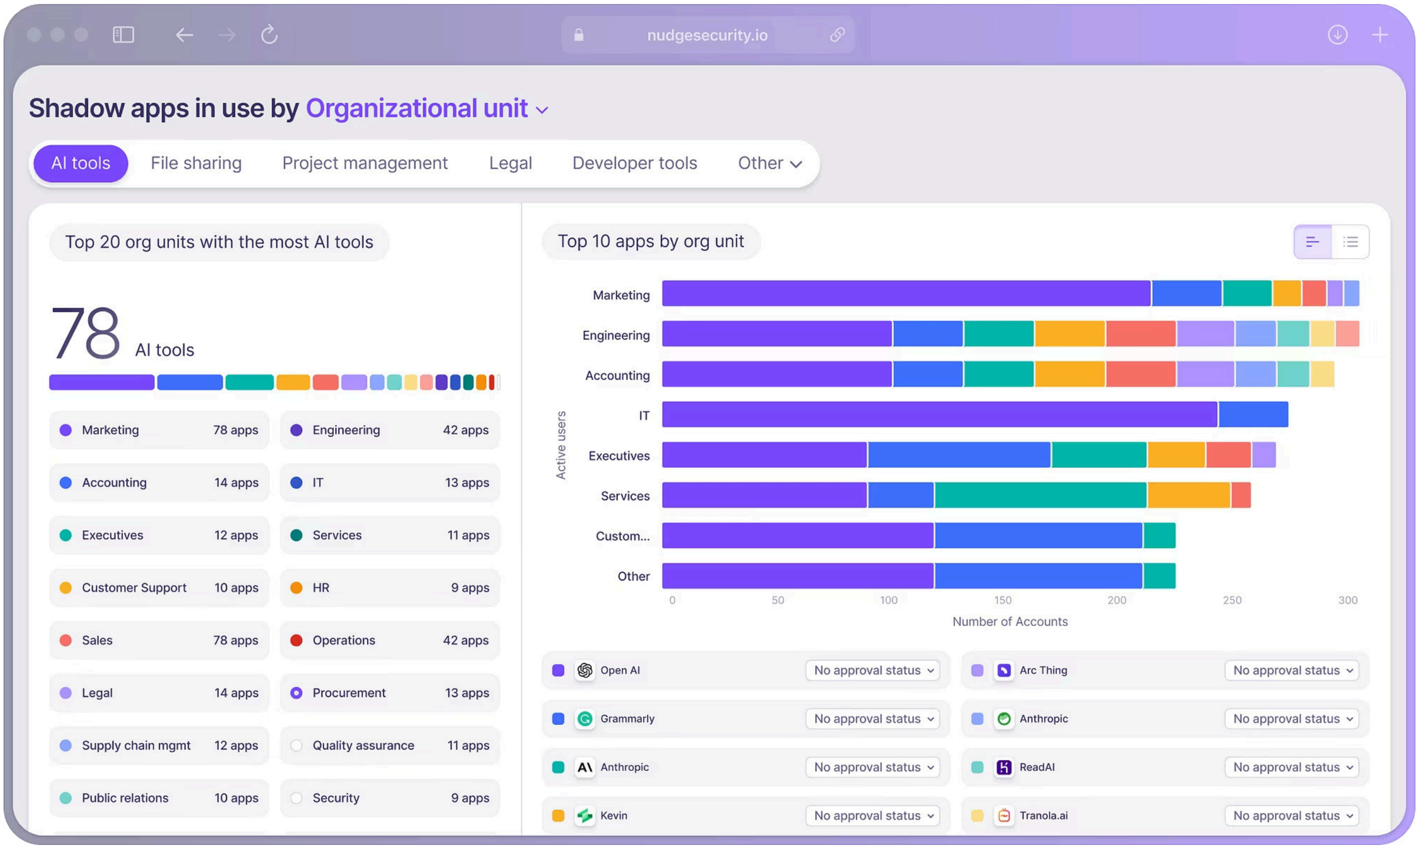

Shadow apps in use by Organizational unit

Top 20 org units with the most AI tools

78 AI tools

Marketing	78 apps	Engineering	42 apps
Accounting	14 apps	IT	13 apps
Executives	12 apps	Services	11 apps
Customer Support	10 apps	HR	9 apps
Sales	78 apps	Operations	42 apps
Legal	14 apps	Procurement	13 apps
Supply chain mgmt	12 apps	Quality assurance	11 apps
Public relations	10 apps	Security	9 apps

Top 10 apps by org unit

Open AI	No approval status	Arc Thing	No approval status
Grammarly	No approval status	Anthropic	No approval status
Anthropic	No approval status	ReadAI	No approval status
Kevin	No approval status	Tranola.ai	No approval status

Nudge Security Differentiators

Nudge Security's discovery capabilities serve as the foundation for a complete SaaS and AI security solution that helps organizations manage their entire SaaS attack surface, from discovery through governance with automated guardrails that guide your workforce to use SaaS and AI in safe, compliant ways.

Comprehensive coverage

Nudge Security's layered approach combines email analysis, API integrations, SSO connections, and a browser extension to provide the broadest possible visibility into AI use.

Automated discovery

Nudge Security continuously identifies new AI apps and integrations as they are introduced, with minimal setup or maintenance, and requiring no prior knowledge of an apps' existence.

Historical insight

Nudge Security uncovers AI accounts created in the past, giving you immediate visibility into your entire AI footprint, even apps introduced before you started using Nudge.

Rich context

Beyond just discovering AI apps, Nudge Security provides detailed information about usage patterns, user roles, authentication methods, and security configurations.

Scalability

Nudge Security doesn't require complex infrastructure, endpoint agents, or network monitoring tools, making it ideal for modern, distributed workforces.

Start a free 14-day trial today

> "Having Nudge has significantly brought peace of mind because I don't have to go looking for a needle in a haystack anymore. This has been my dream that I've been looking for for a long time, for years."
>
> Leo C, IT Team member at GLAAD