The Practitioner's Guide to Conducting an AI Risk Assessment

While most orgs have moved from panic to practicality when it comes to AI use, there are nuances to risk mitigation for GenAI versus other technologies. Learn how to take a practical approach to assessing AI risks.

AI is no longer a futuristic concept; it's a present-day reality transforming how businesses operate. For security and IT leaders, understanding and managing the risks associated with AI implementation is crucial to safeguarding organizational data and reputation. This guide provides a structured approach to running an AI risk assessment, ensuring safe and compliant adoption of AI tools in your enterprise.

Exploring the AI landscape

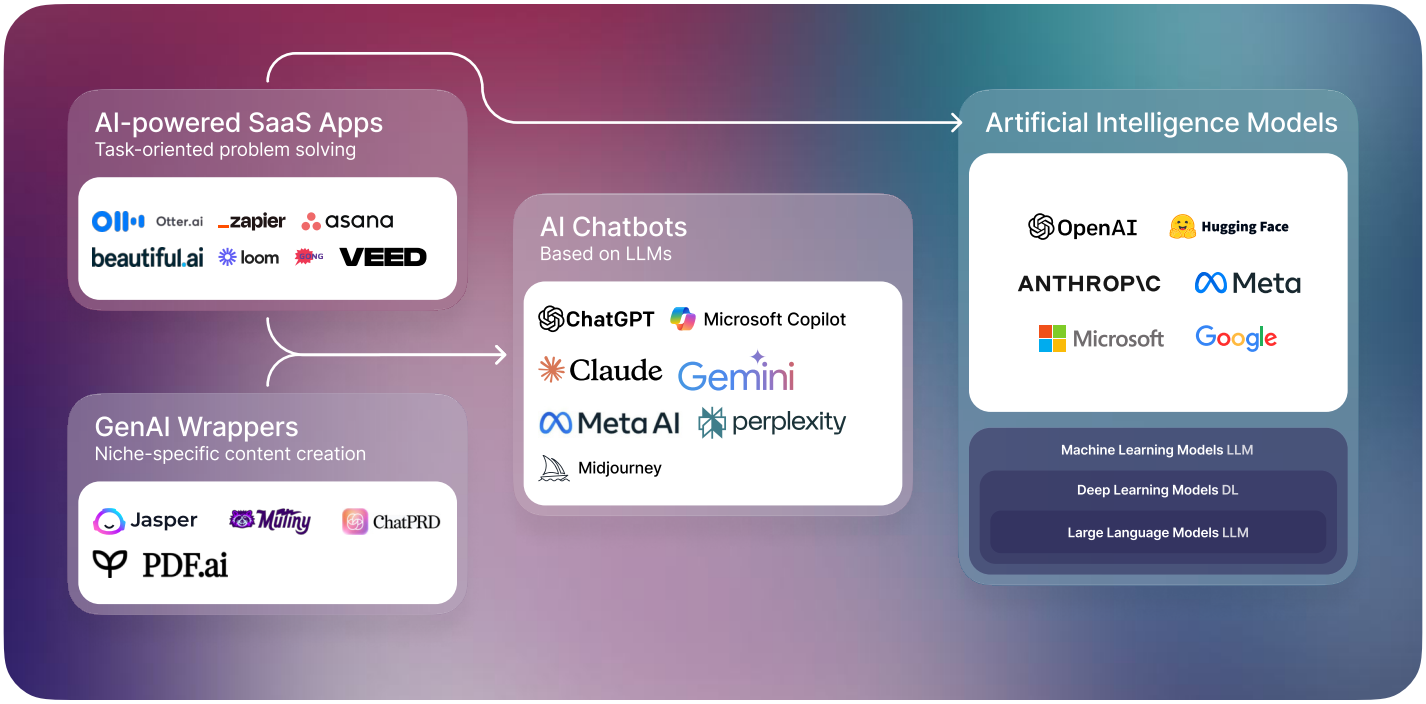

The foundation of AI starts with machine learning, deep learning, and large language learning models (LLMs). The most common models include ChatGPT from OpenAI, Claude from Anthropic, and Llama from Meta. Many of these companies have created a chatbot interface with the models; this is what the general public knows as ChatGPT, Claude, and Meta AI, respectively.

GenAI apps don't stop there. A flood of startups are seizing on the demand for AI by building purpose-built solutions on top of these AI models' APIs. These GenAI "wrapper" apps aim to reduce the learning curve of prompt engineering with a user-friendly UI designed for specific use cases and outcomes. Since they don't require a lot of heavy infrastructural development, GenAI wrappers can be launched quickly and easily as a weekend side project, which may suggest that rigorous security controls are not properly in place.

Finally, there are "AI-powered" SaaS apps: the multitudes of SaaS providers that want to capitalize on the novelty of AI, boost top line revenue, and stay ahead of the competition by embedding AI-powered capabilities in their offerings. "AI-powered" could mean anything from using one of the common LLMs to surface documentation faster to actively delivering suggestions, results, and value within the product.

The bottom line: the AI landscape is vast and growing exponentially faster. In fact, AI growth trends from Nudge Security show that the number of unique GenAI tools has roughly doubled each quarter starting in 2023. It's critical to keep up with the pace that GenAI tools are created and used by your employees and their SaaS tools.

Data from Nudge Security shows that the number of unique GenAI tools has roughly doubled each quarter since 2023.

5 key factors of AI risk assessment

1. Discovery: What AI tools do our employees use?

The first step in any AI risk assessment is identifying all AI-related accounts, users, and applications within your organization. This process involves cataloging not only known GenAI tools in use, but also uncovering new, niche tools that may have slipped under the radar, and any AI-powered SaaS apps. There are five ways to discover what GenAI tools are being used at your organization, providing various levels of visibility.

- **Manual project cataloging:** Teams that want to take an organized, business-aligned approach to exploring and adopting AI often start here. Manual inventory inevitably gets outdated quickly, especially without owners to keep the list updated regularly as business needs change.
- **Network monitoring:** Inspecting traffic details can provide a quick baseline of commonly known GenAI tools, like ChatGPT. These solutions rely on a user-generated database of information to identify specific apps, meaning you have to tell the tool which GenAI apps to look for, rather than the tool discovering new ones for you. While it's a better place to start than a manual inventory, you could still have blind spots when it comes to new, niche GenAI tools—like those that just launched on Product Hunt over the weekend.
- **App integrations:** Analyzing app integrations takes a very narrow view of what GenAI tools are used and connected to core work apps. If you are only concerned with AI risk to critical resources within the organization, this approach may provide a deep view into the connected ecosystem of apps your workforce uses. The downside to this approach, though, is that it is a manual process and may miss any experimentation happening outside of GenAI apps deemed business critical.
- **Agents (e.g. browser extensions, endpoint protection):** While agents can help uncover when and where GenAI tools are used, they may be limited in their visibility of net-new GenAI apps. They also collect way more data than you can put into action. With this approach, it's especially important to be mindful of privacy concerns and ensure proper data handling procedures are in place.
- **Machine-generated email discovery:** Every new account sign-up typically triggers a machine-generated email sent to the email address associated with the new account. Taking an identity-first approach, email discovery covers many of the gaps the other solutions pose from automatically recognizing new, niche GenAI tools to discovering AI use even when the person authenticates with simply a username and password.

2. Trust: How do AI vendors handle security and privacy?

Given the pace at which GenAI tools have entered the market (many without security programs), it's vital to determine the security posture of GenAI tools and ensure they align with your organization's security standards. (This is especially concerning given that 90% of the 1,000+ GenAI vendors Nudge Security has automatically discovered and catalogued have fewer than 50 employees.)

90% of the 1,000+ GenAI vendors Nudge Security has discovered and catalogued have fewer than 50 employees.

When reviewing GenAI vendors, many of the same questions that you would consider for other vendors apply:

- What compliance certifications does the vendor hold?
- Is there evidence of a mature security program?
- How do they manage data privacy?
- Have they experienced any recent data breaches?
- Where is the company headquartered?
- Where is customer data stored and processed?

However, for GenAI apps it's also important to understand if/how you can prevent your data from being used in training models which introduces the risk that it could be surfaced in response to prompts from users outside of your business.

Consider asking these questions to better understand this risk:

1. What types of data are collected from user interactions, and how is it stored?
2. Can data deletion or retention policies be customized for compliance with regulations like GDPR, CCPA, or HIPAA?
3. Can customers opt out of having their data used for training? If so, how, and is that option only available in specific pricing tiers? What mechanisms does the vendor have in place to ensure compliance with customer opt-out requests?
4. Are there contractual agreements (e.g., Data Processing Agreements, SLAs) that explicitly prevent customer data from being used in model training?
5. Does the vendor provide on-premise or private cloud deployment options to ensure data remains within a controlled environment?
6. Does the vendor offer enterprise or dedicated AI models that do not use shared training data?
7. Are there any publicly available disclosures about your AI model training practices?

And, given the complexities of how data is handled in AI tools, questions around data locality and data processing will likely require a closer review than for other types of tools.

By asking these questions, you can better evaluate the AI provider's data policies and determine the level of control you have over your sensitive information.

While security questionnaires can cover some of these questions, conducting these reviews can be time intensive and impede workforce productivity if it is a requirement before using every tool. It can be helpful to steer employees towards already vetted and approved GenAI tools rather than continuing a never-ending stream of GenAI vendor security reviews.

Note: Nudge Security provides free, publicly available security profiles for thousands of SaaS tools, including an expanding list of GenAI tools.

Nudge Security provides free, publicly available security profiles for thousands of SaaS tools, including an expanding list of GenAI tools.

3. Integrations: What GenAI tools are directly connected to core business apps?

GenAI tools often connect to other systems within your organization, creating points where data leaks could happen if not properly managed. A detailed integrations review helps map out these connections and assess their security implications. Key considerations include:

- What systems are connected to the GenAI tools?
- Are they connected via API keys, webhooks, direct integrations, or via OAuth grants?
- What are the scopes for OAuth grants used to connect GenAI tools to other apps?
- What data is being shared with these connected GenAI tools?
- How are the identities managing the integrations protected?
- What guardrails are in place to prevent sensitive data from being shared?

Discovering integrations with AI tools is not always straightforward. A good place to start is to review OAuth grants in your IdP (Microsoft 365, Google Workspace) to look for OAuth grants that enable AI tools to access Google Drive, SharePoint, or other data repositories. This blog post covers more details.

Beyond that, you'll also want to look for OAuth grants and API integrations that connect AI tools to your other systems, particularly those that handle sensitive data like finance systems, HR tools, your CRM, etc. Depending on the logging and API options available for these tools, you may be able to forward events related to OAuth grants and API connections to your SIEM or SOAR. Or, if you are using an SSPM solution, you may be able to get details on integrations for the apps that are managed within your SSPM. If neither of these options are in place, then you will likely need to log in to each app to review the list of OAuth grants and API integrations manually.

4. Supply chain: What SaaS apps use AI under the hood?

SaaS vendors have been launching AI-powered functionality at an accelerated pace since ChatGPT went viral in December 2022, so it is critical for third-party risk management teams to stay on top of which vendors are adding AI to their sub-processor list and supply chain.

To manage third and fourth party risk, this review should regularly investigate:

- What new AI has been added to the software supply chain or sub-processor list of their third-party vendors?
- Do these new GenAI tools align with my company's risk appetite for AI?
- How are these SaaS vendors handling data privacy with the new additions of AI in their pipeline?
- How do these SaaS vendors intend to respond to any data breaches that their GenAI tools experience?

5. Communication: How do users know what safe, acceptable AI use looks like?

Employees want more AI education. In fact, in a recent EY survey, 81% say they would feel more comfortable about using AI if best practices on responsible AI use were routinely shared. IT and security leaders have an opportunity and responsibility to ensure that employees are aware of the organization's acceptable use policy and AI best practices. Regular training sessions, clear communication channels, and accessible support resources can help reinforce these policies.

Ask yourself:

- Have all users been trained on the acceptable use policy?
- What automation can I put in place to ensure my employees are educated on our AI policy as they begin to use GenAI tools?
- Do they know where to find help if they have questions or encounter issues?
- Do they know what approved GenAI tools are available to them?
- Do they know how to request approval for a new GenAI tool or new project involving AI?

81% of employees say they would feel more comfortable about using AI if best practices on responsible AI use were routinely shared.

In Summary

By embracing a structured approach to AI risk assessment, leaders can not only safeguard their data and reputation but also unlock AI's transformative potential securely. Encouraging a culture of vigilance and continuous improvement positions your organization at the forefront of innovation while maintaining robust security protocols.

How Nudge Security can help

Nudge Security has discovered over 1,000 unique GenAI tools in customer environments to date, and provides a scalable approach to AI governance. With Nudge Security, you can:

- Immediately discover every GenAI app and account ever introduced to your org, including tools you've never heard of and accounts created before Nudge was even deployed.
- Review security profiles for every app to assess risks and speed up vendor security reviews.
- Prompt employees to review and acknowledge your GenAI acceptable use policy as soon as they create a GenAI account.
- Publish a directory of approved apps to guide employees towards vetted options.
- See where AI tools are integrated with other tools and where AI is in the supply chain of other apps.

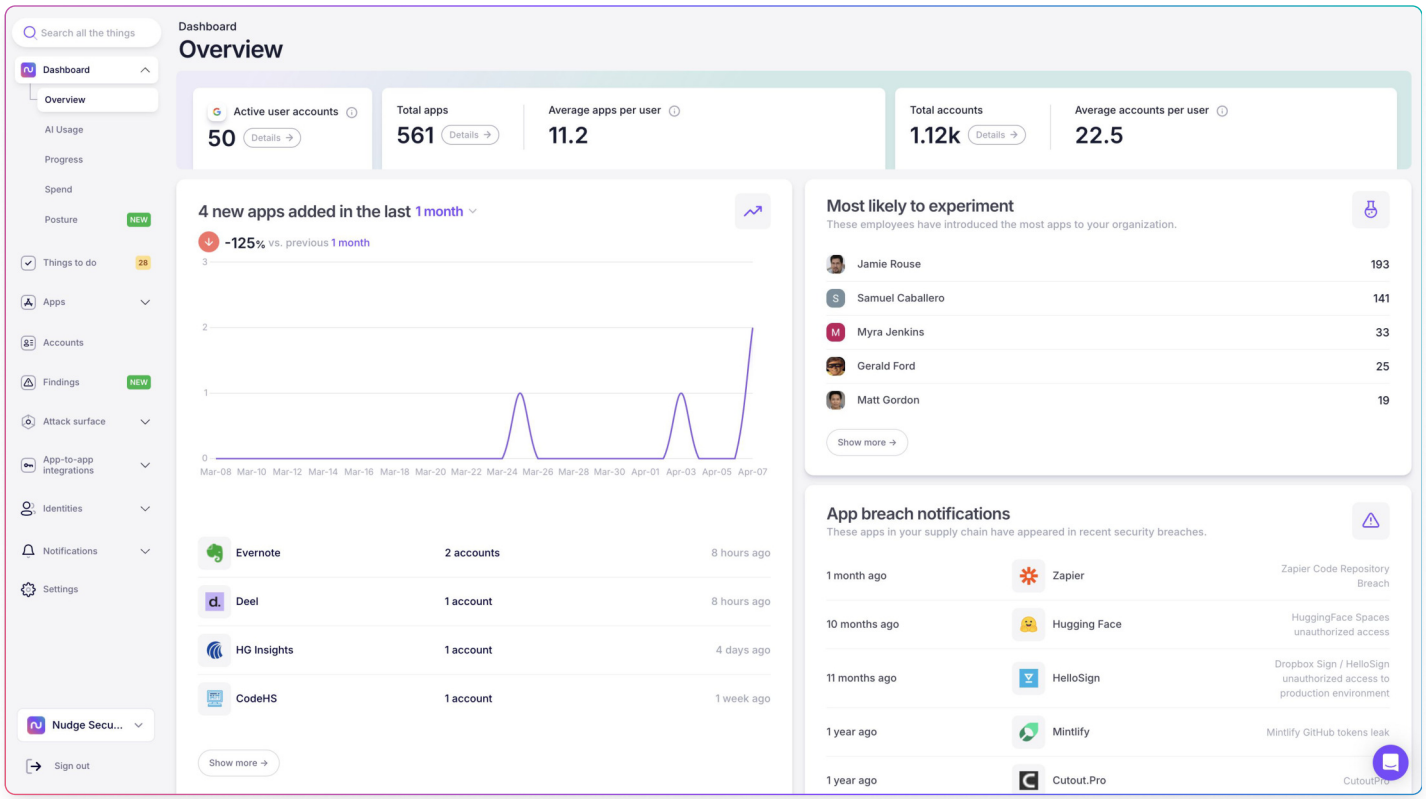

[Learn more about our approach to GenAI security →]

nudgesecurity.com