eBook +

Securing Healthcare Manufacturing in the AI-Driven Threat Landscape

What AI, cloud, and data risks really look like in healthcare manufacturing—and how to stay secure.

netskope

Table of Contents

Introduction

The changing threat landscape for healthcare manufacturing

Manufacturers in pharmaceuticals, medical devices, and the broader healthcare industry are rapidly adopting AI. In pharma alone, AI investment is forecast to grow from $4bn in 2025 to $25bn by 2030, according to McKinsey.[1]

Healthcare manufacturing has always been a highly technical sector, from its carefully controlled chemistry to its statistical modeling and cleanroom facilities. Now with AI, healthcare manufacturers can innovate even further.

AI helps businesses to optimize their production environments, improve efficiency, and accelerate discovery. This new technology can power everything from quality control and predictive maintenance, to process improvements and the creation of regulatory documentation.

But while AI supports safer products, faster innovation, and more effective workflows, it also introduces new points of digital exposure, with significant security consequences.

The widening attack surface

As healthcare organizations accelerate digital transformation through connected medical devices and AI experimentation, their attack surfaces widen, exposing them to unprecedented risks.

In this hostile landscape, legacy security models are failing to protect highly sensitive PHI and critical clinical workflows that form the backbone of patient care. New and emerging risks include:

- Intellectual property theft
- Data leakage across distributed teams and suppliers
- Sensitive files moving through cloud apps and unmanaged devices
- Unsecured AI use exposing proprietary data and regulated content
- Growing access and identity complexity across converging IT/OT environments and suppliers

With manufacturers automating more of their operations, digitizing key workflows, and collaborating more closely with external partners, the boundaries of the production environment start to blur. Data, applications, equipment, and access rights all begin to sprawl.

The result is a more dynamic and efficient manufacturing ecosystem, but it's also one that is harder to keep tightly controlled.

As healthcare organizations accelerate digital transformation through connected medical devices and AI experimentation, their attack surfaces widen, exposing them to unprecedented risks.

The AI impact

Based on exclusive data from our Netskope One platform, this eBook explores one of the key challenges facing the healthcare manufacturing sector today: the impact of AI on the threat landscape. It highlights the balance that organizations are aiming to strike, driving forward rapid innovation while also strengthening data protection, compliance, and risk management controls.

To secure AI, cloud applications, and sensitive data, healthcare manufacturers must move beyond the traditional perimeter-based defense model. Instead, they need to adopt a modern, adaptive framework to security built on zero trust principles, assuming breach by default and prioritizing continuous access verification, granular data protection, and real-time visibility. The rest of this eBook sets out why that's so important.

As generative AI introduced new security risks, embecta, a healthcare manufacturing company, worked with Netskope to modernize its security architecture, replacing traditional network controls with a zero trust approach built for cloud, AI, and remote users.

The big picture: AI in healthcare manufacturing

AI goes mainstream

Generative AI (genAI) tools have been adopted at a remarkable pace in both healthcare and manufacturing contexts. Netskope data shows that 82% of healthcare organizations[2] now use cloud-based genAI apps, while that figure is at 94% among manufacturers.[3]

In both sectors, use of apps that leverage user data for training is even higher (96% in healthcare and 97% in manufacturing). Likewise, the use of apps with embedded genAI features is almost universal (98% in healthcare and 96% in manufacturing).

Healthcare manufacturers' choice of genAI applications largely mirror global adoption patterns. As of May 2026, ChatGPT is the most popular app, with 98% in healthcare ([4] as of May 2026) while 99% of manufacturing organizations using it, followed by Google Gemini (74%).

As those statistics indicate, genAI is now completely mainstream in manufacturing and healthcare organizations, and deeply embedded in their daily work. Far from being an emerging challenge, as it's sometimes described, AI is already here and posing risks today.

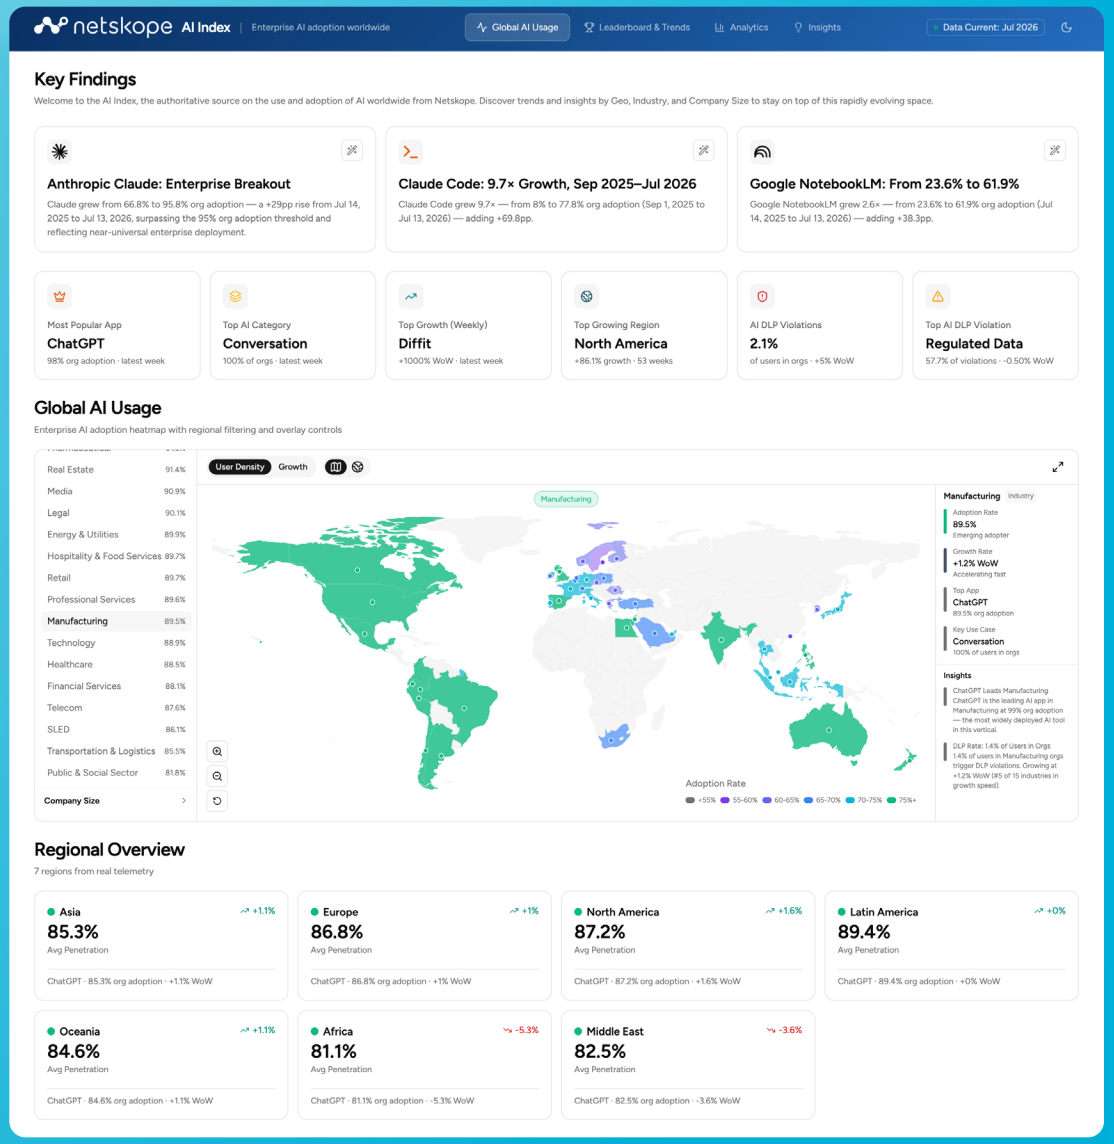

From personal to approved genAI

Since the emergence of genAI in 2022, one constant risk for security teams is the use of personal tools or accounts in the workplace, a practice known as shadow AI. While this remains an issue, the good news is that shadow AI risk appears to be on the wane.

Among manufacturers, personal genAI usage has declined precipitously from 83% in December 2024 to 51% by September 2025. At the same time, the prevalence of corporate-approved genAI tools has risen from 15% to 42% over the same period.

In healthcare, there has been a more dramatic shift, with use of personal genAI applications dropping sharply from 82% to 32%, while adoption of managed genAI solutions has risen significantly from 12% to 56%.[4]

In both sectors the trajectory is the same, moving from individual experimentation toward sanctioned AI platforms with better governance and oversight.

Manufacturing

GenAI usage personal vs. organization account breakdown in the manufacturing sector

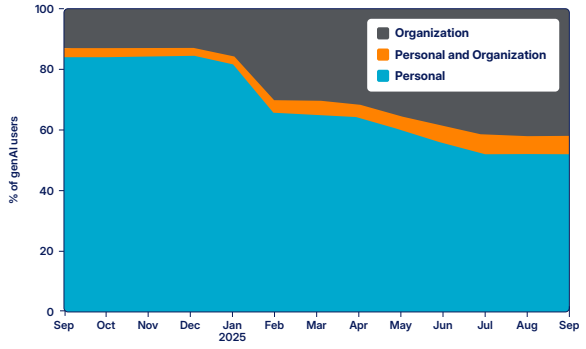

Healthcare

GenAI usage personal vs. organization account breakdown in the healthcare sector

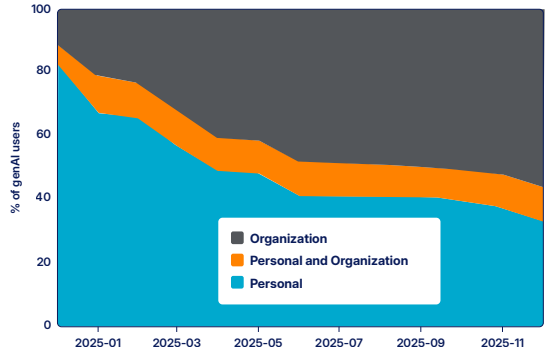

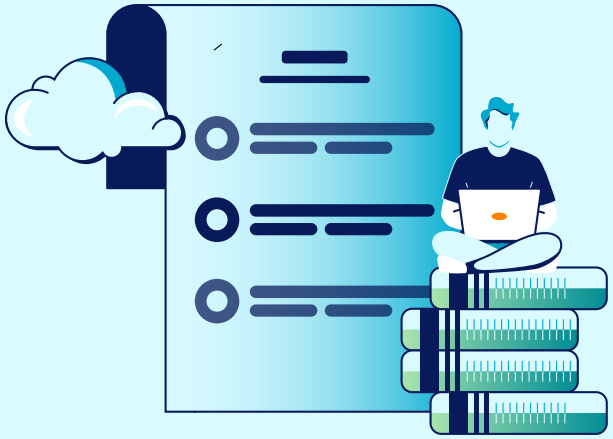

Maturing at pace

Healthcare manufacturing is following the same maturity curve as the wider manufacturing sector, moving from personal experimentation to widespread, sanctioned use. There's already a shift toward embedded AI, APIs, and agentic workflows.

This signals an industry that's no longer testing genAI at the edges but integrating it into core processes. Given the value of speed and innovation in healthcare manufacturing, and the large volumes of data it has to analyze, the attraction of genAI is understandable.

However, deploying AI more widely and embedding it deeper into organizations also creates new points of digital exposure that need to be secured at scale. We'll explore this further in the coming chapters.

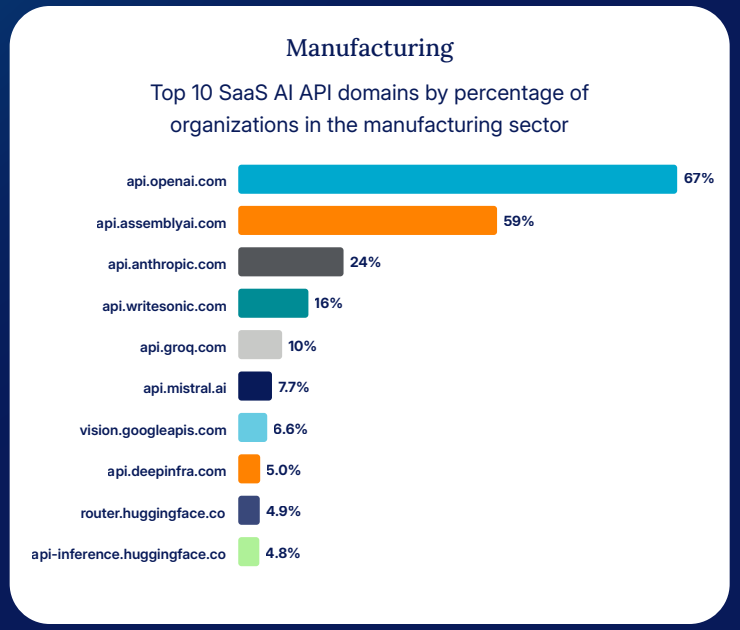

The new security frontier

Even when genAI applications are deployed on premises, the underlying models are often hosted in the cloud, and usually connect through dedicated API endpoints. Among manufacturing companies, 67% now connect to OpenAI's API, while 59% use APIs belonging to AssemblyAI and 24% use Anthropic's.[3]

This approach can help enterprises automate workflows and different elements of the production system, speeding up everyday work, and reducing the administrative workload for human staff. It also creates new pathways for threats to enter the organization.

More fundamentally, this approach takes AI out of browser chats and embeds it into internal tools, agents, and broader systems, with potentially serious consequences. While much of the security focus until recently has been on the threats posed by genAI, practitioners must now ensure they're able to manage the new risks raised by growing use of APIs and agentic tools.

AI risk extends beyond the chatbot. APIs, agents, and integrated workflows create new routes for data exposure and threat activity.

Cloud threats: Malware hiding in trusted apps

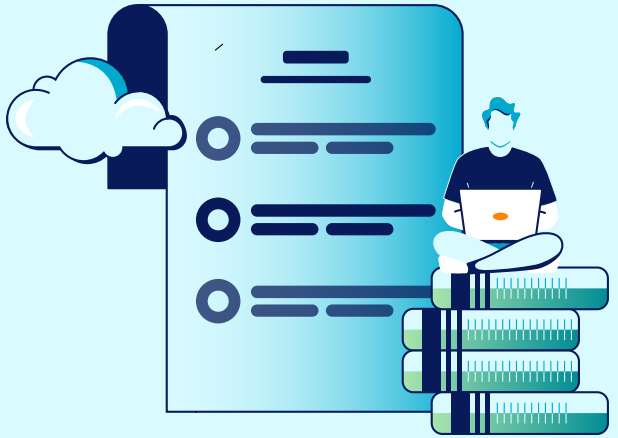

The process of developing new drugs, therapies, devices, and components relies on a complex, often multinational, web of skilled staff and digital systems.

Because of the industry's geographically dispersed workforce and cloud-based workflows, malware is a significant risk. Once successfully inside an organization's network, malware infiltrates an enterprise system to steal sensitive data, disrupt operations, and spread laterally across networks.

Cloud apps as delivery channels

Bad actors regularly use popular cloud applications as their vehicle for distributing malware because they know people are more likely to interact with documents hosted on familiar services. These malicious files blend into normal workflow traffic, bypassing legacy defenses. Just one compromise can propagate malware rapidly across teams and supply chains to disrupt operations or expose important information.

How big a challenge is malware delivered via common cloud platforms? In the manufacturing sector, 18% of organizations see malware downloads from Microsoft OneDrive each month, while 14% get them via GitHub and 11% from Google Drive.[3] In healthcare, GitHub is the most abused service, with 13% of organizations seeing malware downloads from there, although OneDrive, Amazon S3, and Google Drive are also exploited by attackers.

While these platforms actively remove malicious content as soon as possible, the brief window before detection can be sufficient for attacks to succeed. This reinforces the need for security teams to have continuous oversight into network usage and data flows.

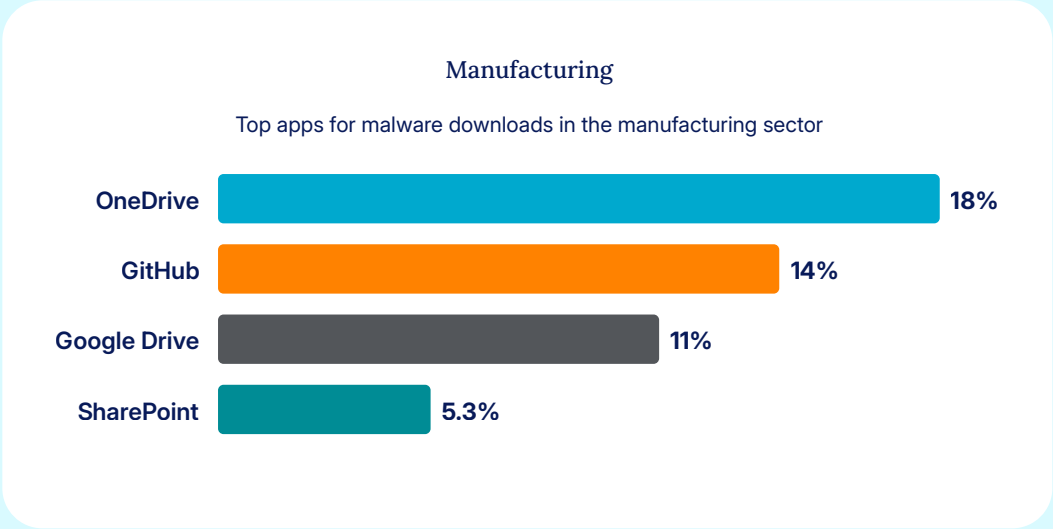

Collaboration risks

At the heart of healthcare manufacturers sit the R&D teams and device engineers, who routinely share key assets relating to product discovery and development. This might include CAD files, designs, product specifications, or compliance documentation.

A single malicious file introduced into a trusted repository, shared folder, or collaborator link can have an outsized impact. Because these environments are assumed to be safe, files often evade heightened scrutiny and move quickly across teams and systems. Once inside, infostealers or ransomware can spread into environments that hold regulated data, production details, or high-value IP.

Perimeter limits

Traditional perimeter-based defenses struggle to handle these new dynamics. They assume that threats arrive from outside, whereas today's digital threats are often embedded within trusted, collaborative workflows and systems.

The challenge for healthcare manufacturers is that the tools they rely on every day are also the most high-risk delivery vectors. This means that enterprises' exposure grows faster than their ability to police it.

The collaboration platforms that accelerate healthcare innovation have also become a primary pathway for ransomware, infostealers, and intellectual property theft.

Exposure risks: How data escapes corporate defenses

Blurring boundaries

GenAI is not the only area where employees are using their own, non-approved accounts at work. More standard cloud applications also present risks as the boundaries between the personal and professional blur.

Virtually every manufacturing enterprise contains workers using personal accounts on apps like Google Drive (in 98% of organizations), LinkedIn (95%), and Microsoft OneDrive (94%). Healthcare shows similar trends, with widespread personal app usage too.[4] While these apps are often used for legitimate business reasons ranging from collaboration to networking, they can also act as unmonitored exit routes for sensitive information.

Data exposure

The type of information exposed through personal accounts varies, depending on whether we look at cloud-based apps like shared drives and webmail, or individual accounts on genAI chatbots. Regardless, it's regulated data, encompassing personal, financial, and healthcare information, that remains most at risk.

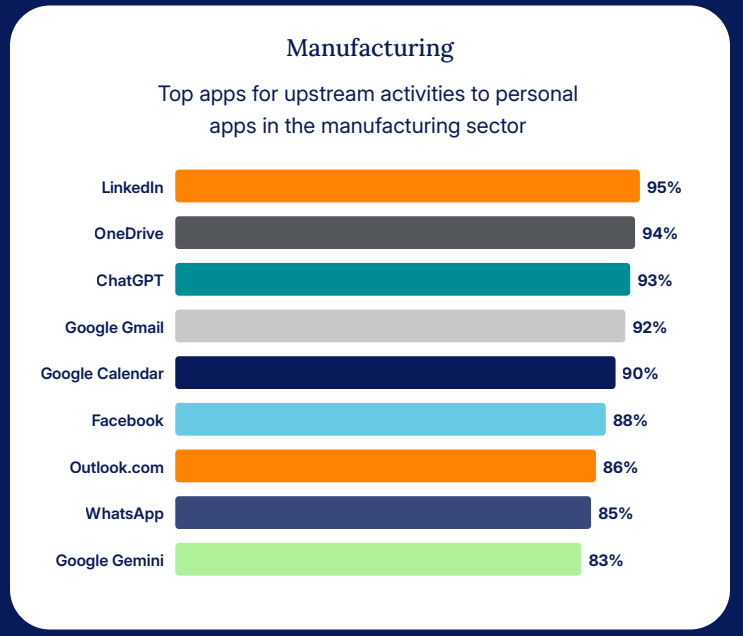

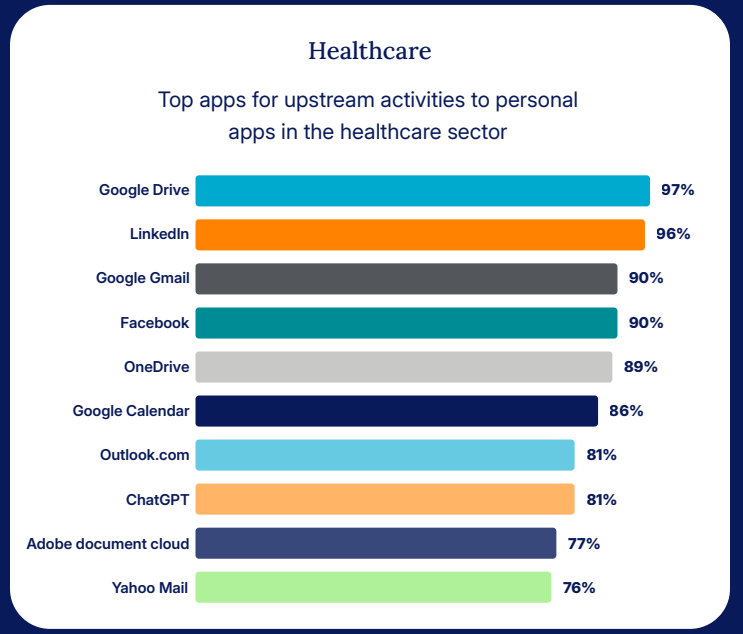

Continued

Analysis of recent incidents at manufacturing organizations shows that regulated data accounts for 41% of policy violations through personal cloud apps.[3] This is followed by exposing intellectual property (32%), passwords and keys (19%), and source code (8%).

Regulated data is also the most frequently lost asset through personal genAI tools, accounting for 29% of all incidents. But it's notable that source code (28%) and passwords and API keys (26%) are much more frequently exposed through personal genAI than through personal cloud apps.

Safeguarding the crown jewels

These findings highlight that regulated data, IP, and source code are the crown jewels of any enterprise, and need constant vigilance to keep them safe.

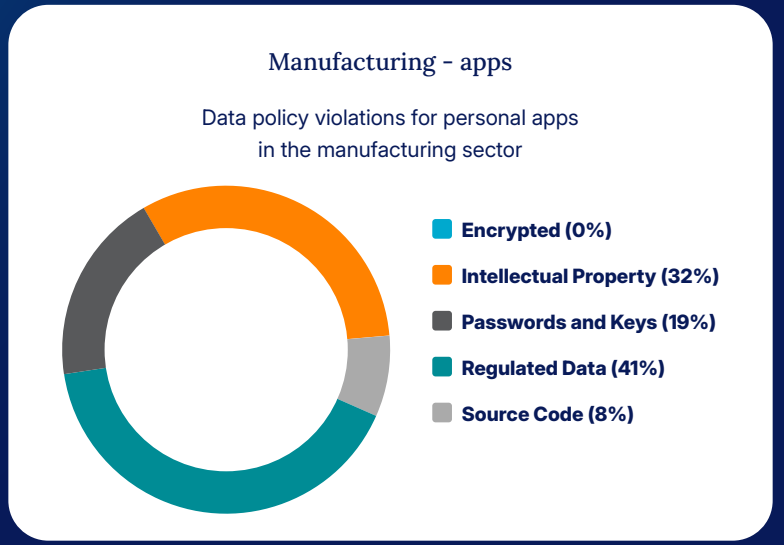

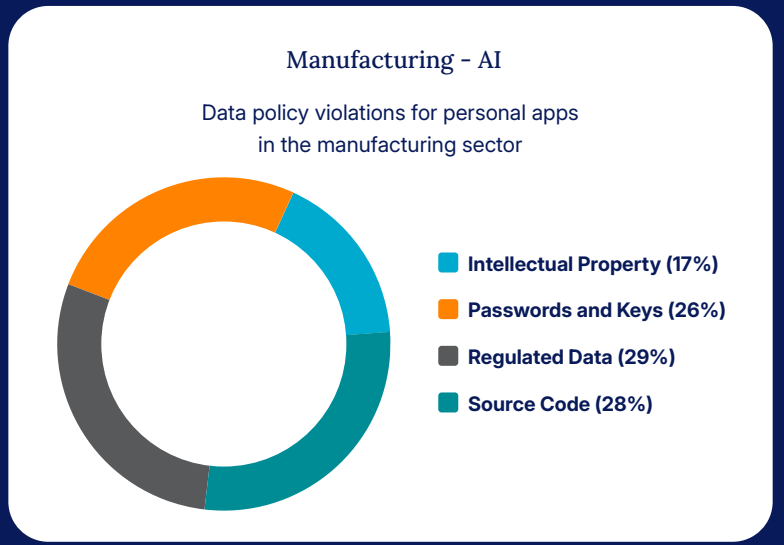

The manufacturer's dilemma

As employees increasingly use unapproved tools to collaborate and work faster, regulated data and high-value IP move outside corporate systems, creating a unique dual risk across two core business pillars:

- **Clinical:** Managing massive volumes of regulated data through validation, quality assurance, and compliance
- **Engineering:** Protecting concentrated IP, including source code and proprietary designs

The convergence of risk

To bridge this gap, healthcare manufacturers require more robust security governance. The next chapter explores how to implement these essential protections.

The business risk extends beyond data loss. Uncontrolled AI use can affect regulatory obligations, product development, and the protection of proprietary knowledge.

Secure healthcare manufacturing's future with the Netskope One platform

The Netskope One platform provides healthcare manufacturing organizations with the security architecture and capabilities they need in the cloud and AI era.

Unified visibility over cloud and web usage

Netskope gives security teams real-time insight into what apps and services are being used, who's using them, and how data flows across environments. Session-level logging, decrypting traffic, and file-level detail. That makes it easier to spot risky behavior, unmanaged services, and unusual data movement.

End-to-end AI security

Netskope secures AI adoption in healthcare manufacturing with visibility and control across AI usage, data flows, and model access. Netskope One AI Security reveals the use of managed and unmanaged generative AI tools, applying consistent semantic data protection policies. The Netskope One AI Gateway secures application-to-LLM interactions, while the Netskope One Agentic Broker controls non-human and agentic workflows. Netskope One AI Red Teaming identifies weaknesses in private models before deployment, and Netskope One AI Guardrails protect runtime interactions from data leakage, prompt injection, and misuse.

Zero trust access and secure connectivity

Traditional VPNs struggle with hybrid work, remote partners, and distributed operations. Netskope One Private Access and the Netskope One SASE architecture utilize the Netskope Zero Trust Engine to continuously verify users and devices while enforcing least-privilege access. Using contextual policy controls, organizations can balance strong security with high user performance and safely integrate AI-powered tools into workflows.

Data protection for sensitive assets

Regulated information must be safeguarded across networks, endpoints, and the cloud. Netskope enforces granular DLP policies, discovers and classifies data with DSPM, and tracks data movement with lineage capabilities.

Threat protection at the edge

Modern threats such as malware, phishing, and ransomware must be detected and blocked in real time. Netskope combines inline threat inspection and behavior analytics to stop advanced threats before they reach key systems.

Simplified security posture and compliance reporting

Healthcare manufacturers benefit from a unified platform and single management console that streamlines policy enforcement, audit reporting, and compliance processes, reducing administrative overhead for security teams.

Governance and controls: How organizations respond

Security teams in healthcare and manufacturing are working to regain control of genAI, cloud apps, and the data flows between them. To do this successfully, they need to move beyond one-off solutions and introduce a consistent governance approach that reflects their ongoing needs.

Targeting controls

Organizations are now applying policies and controls to the personal cloud-based and genAI apps where they see the greatest level of risk. These measures range from blocking uploads, to providing real-time guidance that encourages more informed decisions, to deploying granular data loss prevention (DLP) solutions that can spot sensitive information.

More than one-third of manufacturing organizations (35%) are now protecting Google Drive[3], making it the most controlled app overall. That's followed by personal ChatGPT accounts (controlled by 29% of organizations) and Google Gemini (23%).

Many manufacturers are taking a cautious approach toward certain AI tools, based on their perception of the associated risk level. For example, DeepSeek is blocked by 48% of organizations, while ZeroGPT is blocked by 43%.[3]

The governance gap

It's one thing to invest in these more robust security controls, but all too often they can be applied unevenly. On top of this, existing solutions tend to be designed for traditional enterprise systems rather than today's mix of cloud apps, AI tools, and third-party platforms. The result is a patchwork of protections that still leaves critical blind spots.

These gaps matter more in healthcare manufacturing than in many other industries because of its dual exposure to risk from both regulated clinical data and valuable IP. When governance breaks down, both categories are at risk simultaneously, raising the stakes of even small lapses in control.

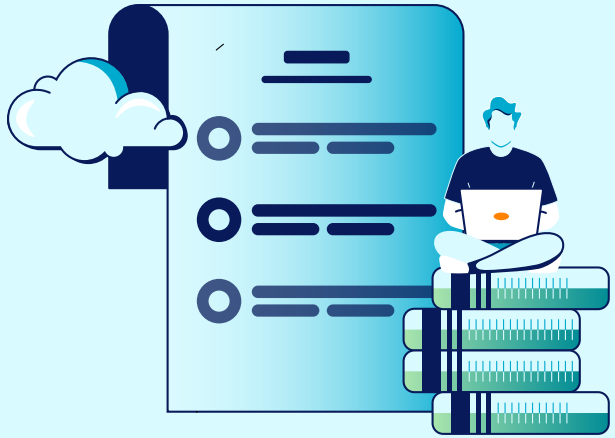

In practice, this dual risk requires governance alignment across often-separated environments:

- Personal and corporate app instances need consistent policies to protect sensitive data.
- Browser and API usage must be governed alongside traditional applications, since they are now primary paths for data movement.
- Cloud storage and embedded genAI tools require the same level of oversight as core systems, especially as they become integral to workflows.
- Critically, these controls must extend beyond internal teams to external suppliers and manufacturing partners who operate inside the same digital ecosystem.

Without this alignment, healthcare manufacturers remain exposed, even as their policy coverage expands. Closing the governance gap is, therefore, less about adding extra rules and more about ensuring that governance reflects how work actually happens. Governance must work for people, systems, and partners, so innovation can continue without compromising security or compliance.

Rethinking governance for the AI era

Security in the AI era requires a fundamental shift in how organizations define and enforce control. Critically, they need to move from app-by-app approval and sanctioning decisions, to implementing robust and consistent data-centric controls.

Traditional governance models were designed for a world where data lived in a small number of approved systems and moved predictably between them. But in a more cloud-based, collaborative, and now AI-driven world, that model is no longer tenable.

In this new environment, governance should be anchored to the data itself rather than to individual tools or platforms. A data-centric approach built on zero trust principles enables consistent policy enforcement wherever data is accessed, whether it's in a corporate SaaS application, a personal cloud account, a web upload, or a genAI prompt. This matters because employees routinely move between approved and unapproved tools to work faster, and data-level controls are the only way to maintain security without slowing them down.

Effective governance in the AI era depends on continuous visibility rather than intermittent oversight. Organizations need a real-time view into how sensitive data is being accessed, shared, and used across systems, so they can respond as risks emerge. This continuous visibility helps turn governance from a reactive exercise into a proactive process that enables innovation and doesn't sacrifice security.

Five takeaways for healthcare manufacturers

Security leaders in healthcare manufacturing organizations face growing challenges as new AI-driven tools and capabilities become commonplace, and are increasingly relied upon for mission-critical work. Throughout this eBook we've provided a snapshot of the threat landscape they face, based on Netskope data and insight.

Risks keep evolving, but here are five key points that security professionals should take out of this overview to help them shape their response.

1. **Expect malware in the tools your teams already trust.**
   The most effective threats today no longer rely on breaking in from the outside. Instead, they arrive through shared drives, collaboration links, and familiar applications. When the files exchanged between colleagues and partners are rarely treated with suspicion, attackers get a fast track into your digital environment.

2. **Assume sensitive data is already flowing through personal and genAI apps.**
   The reality is that employees regularly use personal cloud storage, email, and genAI tools to move faster and get their work done. As a result, regulated data, IP, and source code are already leaving approved systems, often without visibility or oversight. Security teams need tools that can scrutinize every app and block key data when necessary.

3. **AI is no longer a side experiment.**
   What began as ad hoc exploration is now becoming a core capability that organizations must embed into everyday workflows to improve efficiency and competitiveness. Employees increasingly rely on genAI across engineering, R&D, and documentation to accelerate work and improve outcomes. As genAI moves closer to core systems and source materials, its security impact becomes more significant.

4. **Data protection and governance must converge across healthcare and manufacturing needs.**

   Healthcare manufacturers manage both regulated patient data and highly sensitive intellectual property. Historically, these data types have been governed separately, often under different policies, teams, and security controls. However, as both now move through the same cloud platforms, collaboration tools, and genAI-enabled workflows, this separation creates gaps where sensitive information can pass into less-protected environments.

5. **Leaders who act now can turn telemetry into an advantage.**

   The same data organizations already collect to investigate incidents, such as user behavior, file access, and data movement, also provides clear insight into how employees, systems, and partners actually work. Organizations that use this telemetry to align their security with how work actually happens can reduce risk without slowing their innovation.

[1] https://www.mckinsey.com/industries/life-sciences/our-insights/the-synthesis/how-pharma-is-rewriting-the-ai-playbook-perspectives-from-industry-leaders

[2] https://ai-index.netskope.com/ as of April 2026

[3] https://www.netskope.com/resources/threat-labs-reports/threat-labs-report-manufacturing-2025

[4] https://www.netskope.com/resources/threat-labs-reports/threat-labs-report-healthcare-2026

About Netskope

Netskope is a leader in modern security and networking for the cloud and AI era. Built on the Netskope One platform, Netskope unifies secure access, data security, and AI security to give organizations real-time visibility and control across cloud, AI, SaaS, web, and private applications. Powered by NewEdge, Netskope helps customers reduce risk, simplify infrastructure, and eliminate trade-offs between security and performance. Learn more at netskope.com.

netskope