

**WIZ** Google Cloud

# Securing Innovation in the Age of Generative AI





# Responding to the Evolving Threat Landscape

Generative AI and its security implications are a board-level concern at Fortune 500 companies and Global 2000 organizations alike.

Today's security landscape is a cybersecurity race to the finish to keep increasingly powerful AI tools out of the hands of bad actors.

We can't continue to silo security tools based on individual teams, departments, or lines of business. Doing this opens organizations up to threat vectors that exploit AI-based vulnerabilities across the business. Security tools and policies must become holistic, like the cloud on which these AI programs are hosted. A shift from vertical to horizontal security in this new cloud-dominated world is now essential—not just a nice-to-have.

Evolving from vertical siloed solutions to a form factor that serves the needs of every person in the organization requires intense collaboration. Teams, organizations, and industries must share context, language, and policies to safely adopt and innovate with AI in the code.



*We can't continue to silo security tools based on individual teams, departments, or lines of business. Doing this opens organizations up to threat vectors that can exploit AI-based vulnerabilities across the business.*

# Rethinking Code Scanning in the Cloud Era

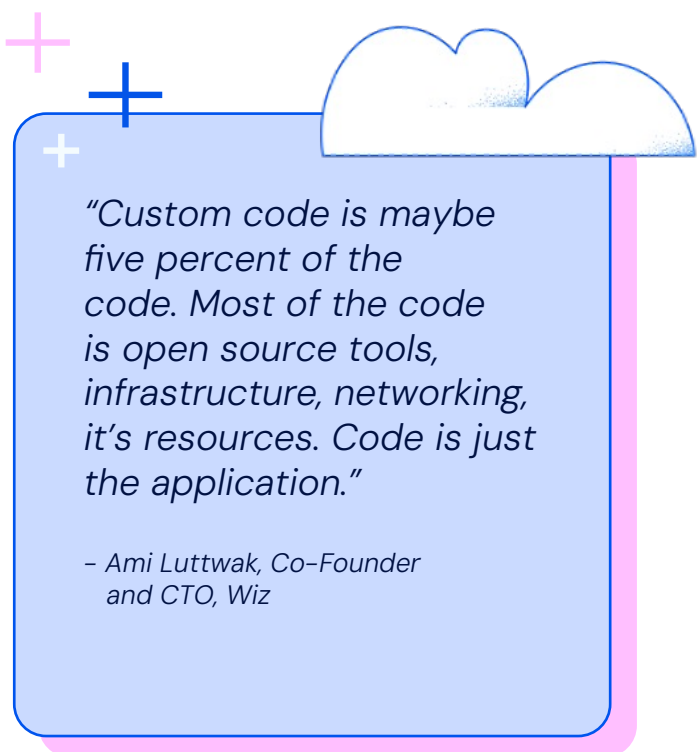
## The Changing Nature of Code

Twenty years ago, code was custom-written by an individual and contained very few instances of what we now call “infrastructure.”

The server’s deployment wasn’t in the code. IT built a rack with server blades. The network wasn’t in the code. A networking person configured routers.

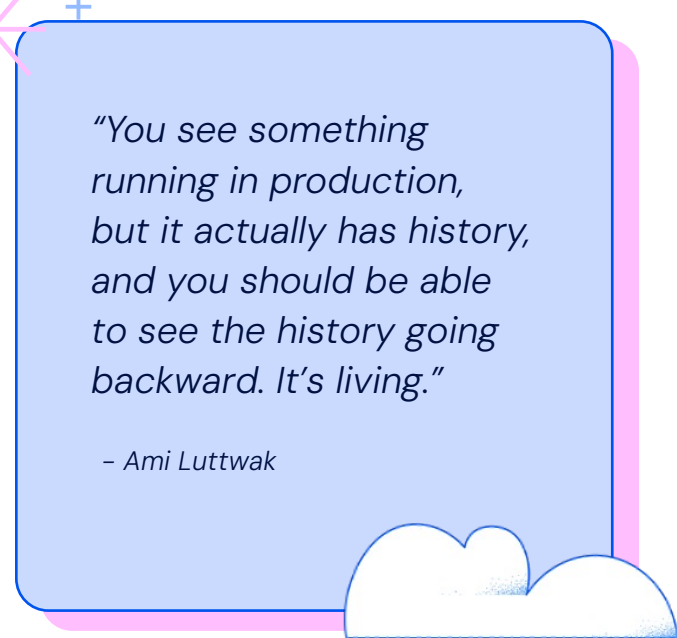
Twenty years ago, code was code. Because of that, it made sense to have a tool for custom logic codes to help find mistakes in memory corruption. This security was simple because other layers, such as application, data, network, and compute, didn’t exist inside the code.

Today’s code and code repositories are light years away from the days of manual IT. Code repositories now contain everything: networking, servers, databases, and applications. This repository is no longer the custom code we once knew. It is a living, breathing application deployed in multiple places simultaneously.



*“Custom code is maybe five percent of the code. Most of the code is open source tools, infrastructure, networking, it’s resources. Code is just the application.”*

– Ami Luttwak, Co-Founder and CTO, Wiz



*“You see something running in production, but it actually has history, and you should be able to see the history going backward. It’s living.”*

– Ami Luttwak

Today, we consider code repositories as an application blueprint. Much as with a building, this blueprint can’t replace the structure, but it can tell us where things are located and where they originated.

However, unlike a building, the application doesn’t exist in only one place. It’s defined somewhere but runs in many places and lives in several versions. The same app can live inside the Terraform code, inside the pipeline, and in production simultaneously.

It no longer makes sense to look at a database in production and fix a problem immediately. Issues are no longer isolated to the area where they appear; problems exist everywhere the application exists.

If we deploy the app daily, it’s more like an organism—constantly changing and adapting to diverse environments.

## How Do We Think About Security for This New Organism?

An application that simultaneously lives in multiple places often creates tension among the teams interacting with it. They all have different points of view, goals, and perspectives on the app in question.

**Developers see the  
app in the code**

**DevOps sees it in all the  
environments**


**Security sees it from the  
point of view of a potential  
attacker**

The fundamental shift in perspective we all must adopt is that every problem is a problem in the code.

If there is an operating system vulnerability, it has a source in the code. The same is true for other alerts that pop up. If everything is in the code, looking at problems only in the cloud no longer makes sense. Development and security teams waste an inordinate amount of time discussing where a problem is located and fixing it in multiple places.



*The fundamental shift in perspective we must adopt is that every problem is a problem in the code.*





## A New Approach to Code Scanning


One thing has become clear as we think about how best to secure work in the cloud. Having different teams is okay, but we must stop buying security tools based on those teams.

*Attackers don't care about our organizational structure.* They're happy when we have siloed tools and teams because, at its heart, this disjointed approach is anti-security. We either need fewer and smaller security teams, or security should be horizontal and integrated.

If the application lives everywhere, then the tools must be able to look everywhere.

However, a shared language and a singular tool don't mean everyone must always see everything. A singular engineer doesn't want to see the entire code landscape, and developers only want to know what's needed to fix the problem.

Security tools that scan the cloud should scan everything. They should be the app's map because the app lives everywhere. If we can implement tools that build context across everything, we can provide the correct information to the right team at the right time.



*"From my perspective, one of the goals of a security tool is to build a shared language that still gives every team the right information at the right time and the right tool."*

– Ami Luttwak

# Balancing Innovation and Governance with Generative AI

How can companies innovate quickly without compromising security?

This is a key question security and governance leaders must address today. Balancing the board's demands for innovating with AI with the need to minimize cloud security risk is a critical challenge.



## The Promise and Peril of AI and the Cloud

The cloud offers benefits that help organizations overcome the challenge of innovating with AI versus cloud security risks. For most organizations, moving to the cloud is mandatory because of the proliferation of AI, but it remains complicated. The question isn't, "Will we move?" It's, "What will we do when we get there?"

Generative AI and large language models (LLMs), including Google, OpenAI, Cohere, Mistral, and Anthropic, have gone from occasional use to ubiquitous. Nearly three-quarters of organizations (72%) now use one or more generative AI functions. The widespread availability of generative AI helps organizations innovate and become more efficient, but it also exposes them to advanced vulnerabilities.

However, despite growing attack vectors brought on by the advent of AI, it remains the single biggest driver of innovation the world has seen in several years.

Here are some of the companies at the forefront of this innovation and breaking new barriers to excellence in their use of AI:

 72%

of organizations now use one or more generative AI functions<sup>1</sup>

<sup>1</sup>The state of AI: How organizations are rewiring to capture value, McKinsey & Company, March 12, 2025.



# Real-World Use Cases

## Example: Cox2M

Cox2M delivers IoT solutions through two flagship products: **GearTrack** and **Kayo**.

**GearTrack** optimizes industrial and supply chain operations by using IoT sensors to monitor high-value assets in transit. By tracking real-time location and condition, it reduces theft, improves efficiency, and enhances inventory management.

**Kayo** empowers small and medium-sized businesses to manage their fleets better. With real-time visibility into fleet performance, Kayo helps business owners monitor driver behavior, track vehicle usage, and optimize operations. By reducing idle time, improving fuel efficiency, and enhancing driver accountability, Kayo allows small business owners to focus more on growing their businesses.

Cox2M processes over **2 million IoT sensor messages per hour**, covering **29,000 miles of asset movement** hourly and analyzing **8 TB of data daily**. This data is critical for understanding traffic patterns, predicting delays, and preventing equipment failures.

However, high costs, long analytics lifecycles, over \$142,000 per data request, and weeks for upgrades hindered experimentation and analytics development.

To address this, Cox2M built a data lake in BigQuery, centralizing its data, and implemented a natural language search engine powered by LLMs. This enabled:

- Intuitive natural language queries
- Real-time follow-up questions for deeper insights
- Automated scheduling of queries and reports

By modernizing its data infrastructure, Cox2M accelerates innovation and delivers actionable insights to optimize supply chains and fleet operations.

These changes significantly improved operational efficiency, resulting in:



**88%**

faster report generation

**75%**



improvement in data structuring

*"A lot of data, as you probably know, means a lot of opportunity."*

*– Stephen Orban, VP of Migrations, ISVs and Marketplace, Google Cloud*

## Example: Oneworld

When planning a trip across all available airlines, a person could encounter 1034 possibilities. Planning quickly becomes unfathomably complex once cost constraints, frequent flyer benefits, and multiple legs or stopovers are layered in.

Oneworld saw an opportunity to create a multimodal ticketing agent to help customers consider all these possibilities.

Its proprietary LLM, built on top of Google Cloud's Vertex AI platform, translates natural language queries into code that matches the structure of its data sets. This approach limits the immense number of flight possibilities. The LLM even weights and scores customers' frequent flyer statuses to narrow the field further.

Oneworld built all of this in six months and has seen a 5x increase in conversions across itinerary purchases.



increase in conversions  
across itinerary purchases

## Example: Major Retailer

In addition to its chain of brick-and-mortar stores, much of this retailer's business has shifted online. One distinctive feature is the sheer size of its catalog. With hundreds of millions of products, the company previously relied on manual metadata tagging and labeling for each item.

The retailer adapted its own LLM to scan product images and automatically add text and vector embeddings to its product pages. This eliminated the need for manual work and resulted in much higher search accuracy for its online customers.

In addition to being much more efficient than manual labeling, the retailer gained:

- More relevant tags
- Enhanced search efficiency, allowing customers to find products faster

The retailer estimates that it is 100 times more productive using generative AI than using people to tag its products.

100x

times more productive using  
generative AI than using  
people to tag its products




## Example: Exabeam

The Exabeam New-Scale Security Operations Platform provides security log management and behavioral analytics to help professionals automate threat detection, investigation, and response (TDIR).

The velocity of threats and the incremental growth of data place tremendous pressure on security operations teams to keep pace. Exabeam Copilot was introduced to help solve this problem. It combines the power of generative AI with exceptional Exabeam user and entity behavioral analytics (UEBA). Built on top of Google Gemini, Exabeam Copilot turns natural language searches in any language into powerful queries on the New-Scale Platform.

These capabilities allow junior security analysts to deliver faster and more precise TDIR without having to understand SQL or the underlying data. In addition to natural language search, Exabeam Copilot offers automated threat summaries, giving responders a glimpse at the potential initial impact of a threat with remediation suggestions. Exabeam Copilot provides tangible insights to decrease the time to respond, leading to better security outcomes.



*The velocity of threats and the incremental growth of data place tremendous pressure on security operations teams to keep pace.*



# Best Practices for Secure AI Development

The increasing use of AI in business has driven new methods and techniques for achieving excellence in AI usage.

It's unclear how this new era will unfold, but the question remains: How do we balance the speed at which our businesses demand innovation with security and governance?

Here are some best practices for AI development and usage to ensure the organization gets the most out of its innovations while remaining protected.

## One Last (Best) Use Case: L'Oreal


L'Oreal, the world's largest cosmetics company, wanted to balance the demand for speed with security and governance. It saw the potential of putting generative AI in the hands of developers, marketers, and business lines without sacrificing governance of its sensitive data or leaking intellectual property.

L'Oreal had its Cloud Center of Excellence team create "Gen AI As a Service." This guardrail set of descriptive generative AI APIs allowed teams to build AI tools and products consistent with L'Oreal's data policies.

The question that plagued L'Oreal—to choose between innovation or governance—became choosing both innovation and governance.

L'Oreal uses tools such as Langchain and Google Gemini to back its APIs. With these tools, developers built chat interfaces, picture and video inquiries, and image generation capabilities to create pages for its cosmetics. This is done with the speed and confidence required to ensure it's not leaking data and sensitive IP.

This solution quickly became a best practice to balance innovation and governance so organizations aren't forced to choose.



## Google Cloud's Secure AI Framework (SAIF)

As AI technology rapidly advances and threats continually evolve, the challenge of protecting AI systems, applications, and users at scale requires that developers have a high-level understanding of AI-specific privacy and security risks in addition to established secure coding best practices.

Google Cloud has since released its **Secure AI Framework**, which offers guidance for building and deploying AI responsibly.

SAIF describes Google Cloud's approach to addressing AI risks—including the security of data, models, infrastructure, and applications involved in building AI—and is aligned with Google Cloud's Responsible AI practices to keep more people safe online. SAIF is designed to help mitigate risks specific to AI systems, such as model exfiltration, data poisoning, injecting malicious inputs through prompt injection, and sensitive data disclosure from training data.



## Building a Collaborative Security Culture

Vertical security means teams deploy tools based on the security organizational structure: vulnerability tools for vulnerability teams, network tools for network teams, and DevSecOps tools for DevSecOps teams.

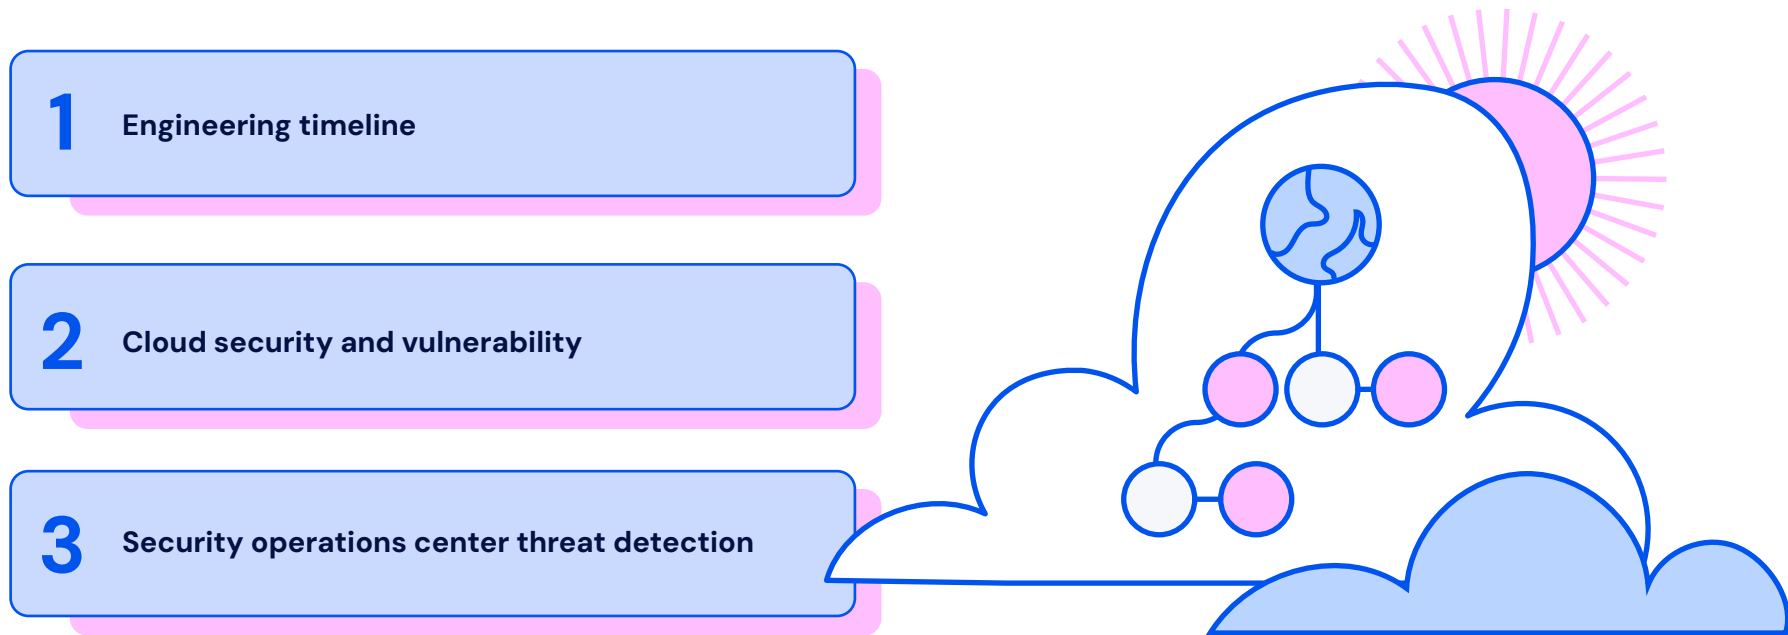
The big problem with this approach is that attackers and engineers don't care about the organizational structure.

Taking a horizontal security approach means we no longer work as a soundproof box from team to team. We empower teams, giving them the correct information at the right time. Horizontal security means we can focus on solving more significant organizational problems, not just providing team-based solutions. With horizontal solutions, teams are finally able to fully communicate with one another. This communication is critical in ensuring the organization remains secure as it accelerates its AI innovations.

## “Shift-Left” Security for AI

“Shift-left” security is essential in an AI-centric world. Shifting left doesn’t mean deploying everything in the pipeline and being grateful when there are no alerts. When considering comprehensive AI security, shifting left is about finding and validating the real problems. This means helping engineers find and fix the problem’s source with the information when and where they need it.

With collaborative security, we must ask ourselves how we work between three workstreams.



Each workstream has its tools, terms, revamps, and timelines. But there is a way for us to work together.

The answer isn’t simply to “shift-left.” We can’t just deploy parallel tools because that’s what we did before the cloud. The cloud and AI change everything, and we must change along with them.




# The Coalition for Secure AI: A Collaborative Approach

Wiz and Google Cloud believe that security and governance are challenges the entire industry must approach together.

Many standards and governing bodies have attempted to develop security frameworks and standards, including:

- International Standard Organization (ISO)
- National Institute of Standards and Technology (NIST)
- Payment Card Industry (PCI)
- General Data Protection Regulation (GDPR)
- Statement on Standards for Attestation Engagements (SSAE) 16 Certified
- Sarbanes-Oxley Compliance
- Health Insurance Portability and Accountability Act (HIPAA)



These frameworks provide a common language and approach to securing intellectual property in digital scenarios. This is why Google Cloud and Wiz have partnered with other cloud providers, hyperscalers, and LLM providers to develop an entirely new entity.

## The Coalition for Secure AI (CoSAI)

Its mission is to tackle the fragmentation in AI security by fostering a collaborative ecosystem. Through CoSAI, we aim to bring together experts from various fields to establish and share best practices, tools, and methodologies for secure AI development and deployment.


The coalition comprises several partners, including Google Cloud, Wiz, NVIDIA, Microsoft, Amazon, Cisco, Anthropic, Intel, and OpenAI. The [complete list of sponsors is available here](#).

There are three founding work streams that CoSAI is focused on:

**1 Software supply chain security for AI systems**

**2 Preparing defenders for a changing cybersecurity landscape**

**3 AI security and privacy governance**



*"Our aim is to build and give the community a set of standards for customers who don't want to trade off speed and innovation for good security and governance."*

– Stephen Orban

Our goal is to find partners throughout the industry, so [reach out](#) to learn more or if there's interest in participating. Together, we can continue discussing how these work streams evolve, what other work streams we can take on, and how we can all work together to arrive at this balance of innovation and security.

CoSAI aims for everyone in the industry to innovate quickly without trading off security and governance. It's not a trade-off we should all have to make. We shouldn't have to choose if we approach AI and innovation in a way that leverages best practices from around the industry.



## Accelerate Innovation Securely

Wiz and Google Cloud ensure organizations keep pace with innovation without sacrificing security.

We utilize the tools that scan for problems everywhere to eliminate silos and bring teams together with a shared language.

Together, we are working to make the future of AI safer and better for everyone.

**See how Wiz and Google Cloud help  
organizations innovate safely.**

