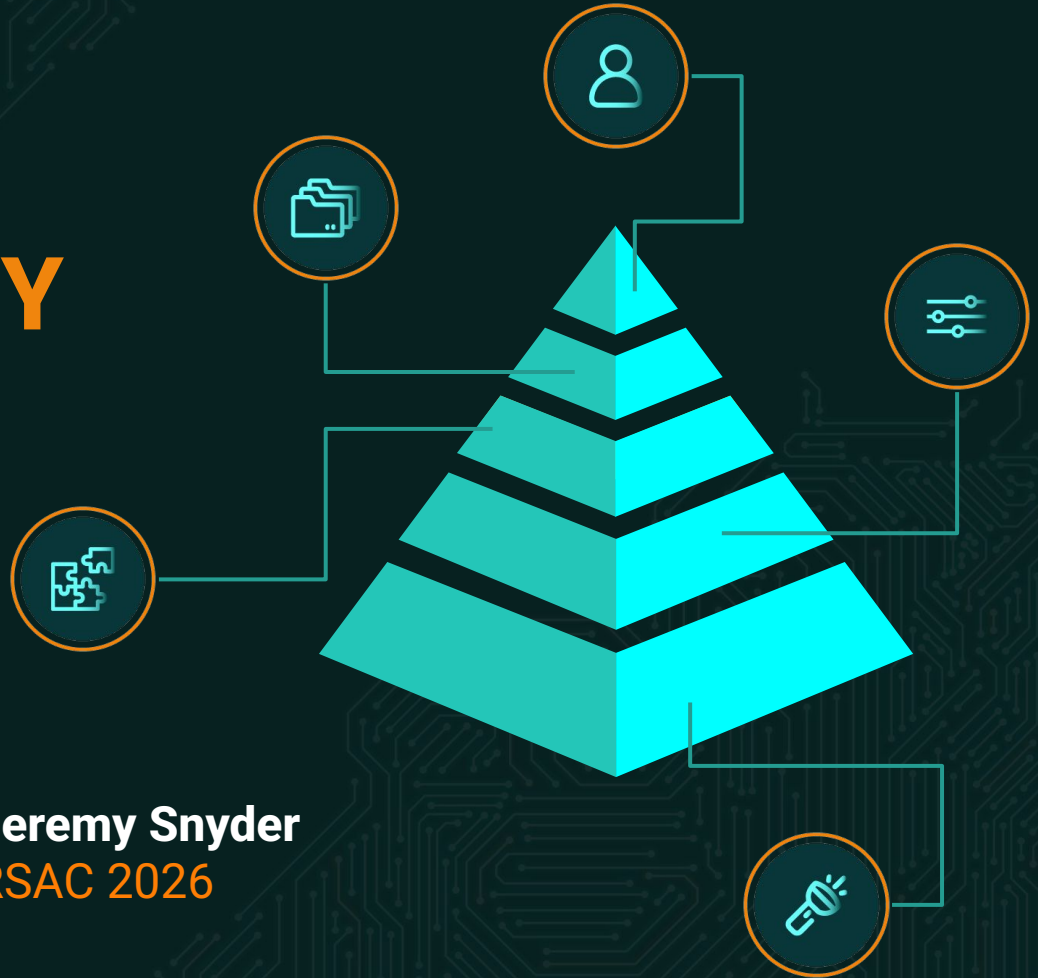


THE 5 LAYERS OF AI SECURITY

What Security Teams
Are Missing



Jeremy Snyder
RSAC 2026



FireTail



Riley Priddle
Co-Founder & CTO



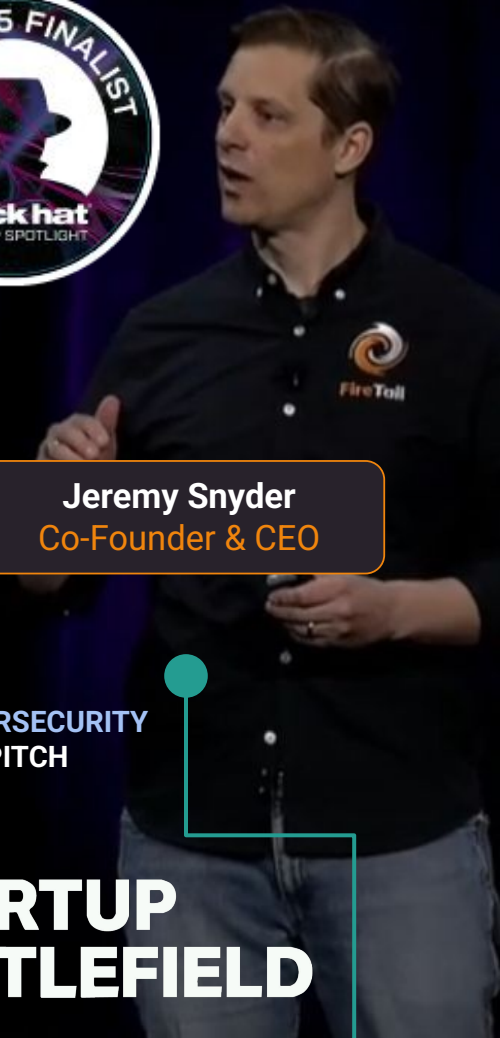
Timo Rüppell
VP of Product



Mikko Hyppönen
Advisor



Jeremy Snyder
Co-Founder & CEO



**Enabling AI adoption is the
biggest challenge facing
security teams today .**



Enabling **secure** AI adoption is
the biggest challenge facing
security teams today .

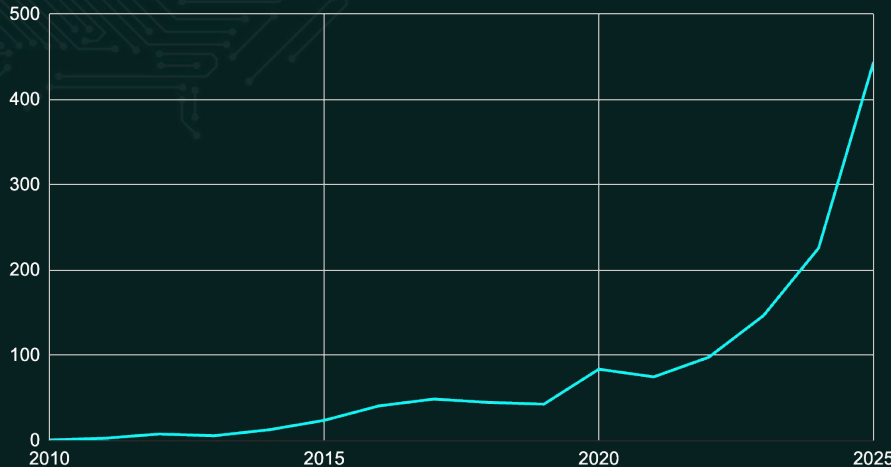


AI Security Concerns Are No Longer Theoretical

- **High-profile incidents** - Some of the world's largest enterprises have had AI-related breaches, leaks and vulnerabilities exposed.



AI Security Incidents



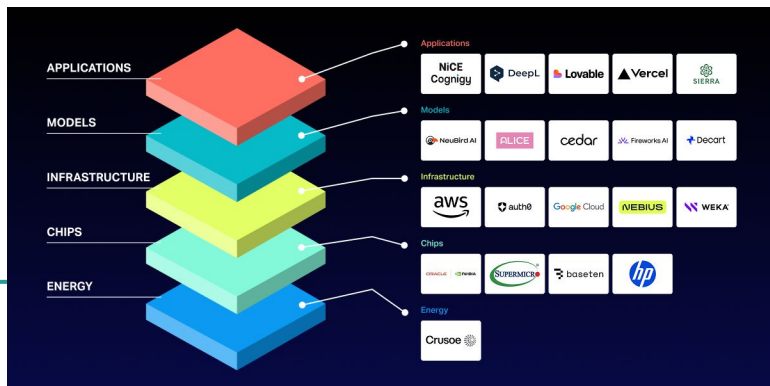
- **Evolving threats** - "By 2027, more than 40% of AI-related data breaches will be caused by the improper use of generative AI (GenAI) across borders."

Gartner

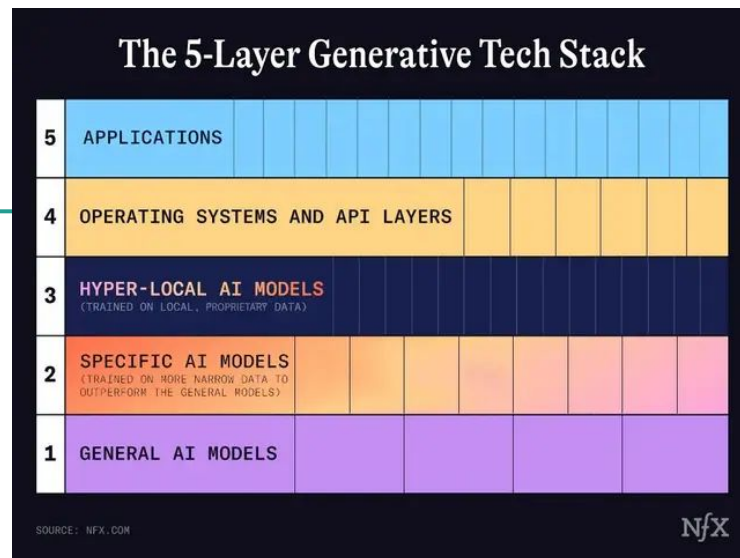
Security in the Age of AI .



Inspiration.



“AI is a 5-layer cake.”
Jensen Huang, Nvidia



“As a founder, you have to decide which layer or layers you want to include in your product.”
NFX

Sources:

NFX, <https://www.nfx.com/post/generative-ai-tech-5-layers>

HumanX, <https://businesschief.com/news/why-jensen-huang-s-ai-strategy-is-built-as-a-five-layer-cake>



Echoes of the past.

Lessons from the cloud

“The reason a startup can build with frontier models today is the same reason a startup could launch in 2006.”

**Werner Vogels,
Amazon**



Werner Vogels  · 1st
VP & CTO at Amazon.com
6h · 

I almost didn't take Amazon's call. It's an online bookstore. How hard could their scaling be?

I did visit to give a talk, and when I took a glance in their (technology) kitchen I was totally blown away. Every distributed systems problem my colleagues and I had been theorizing about at Cornell, from fault tolerance to consistency and to availability at scale, Amazon was wrestling with them live, in production, every day. At a scale orders of magnitude larger than I had ever seen before. And more importantly, no commercial software worked at this scale, so everything was home-grown. Real customers. Real outcome. I wanted to be part of that.

What grew out of that became AWS. And from the beginning, the roadmap was written by customers. A customer couldn't afford infrastructure upfront, so we built EC2. A customer said S3's consistency was breaking their systems, so we fixed it, for everyone, overnight. A customer said the hypervisor overhead was an unbreakable ceiling, so we built Nitro and eliminated it. Every time a customer handed us a task that seemed impossible, we invented our way out of it..

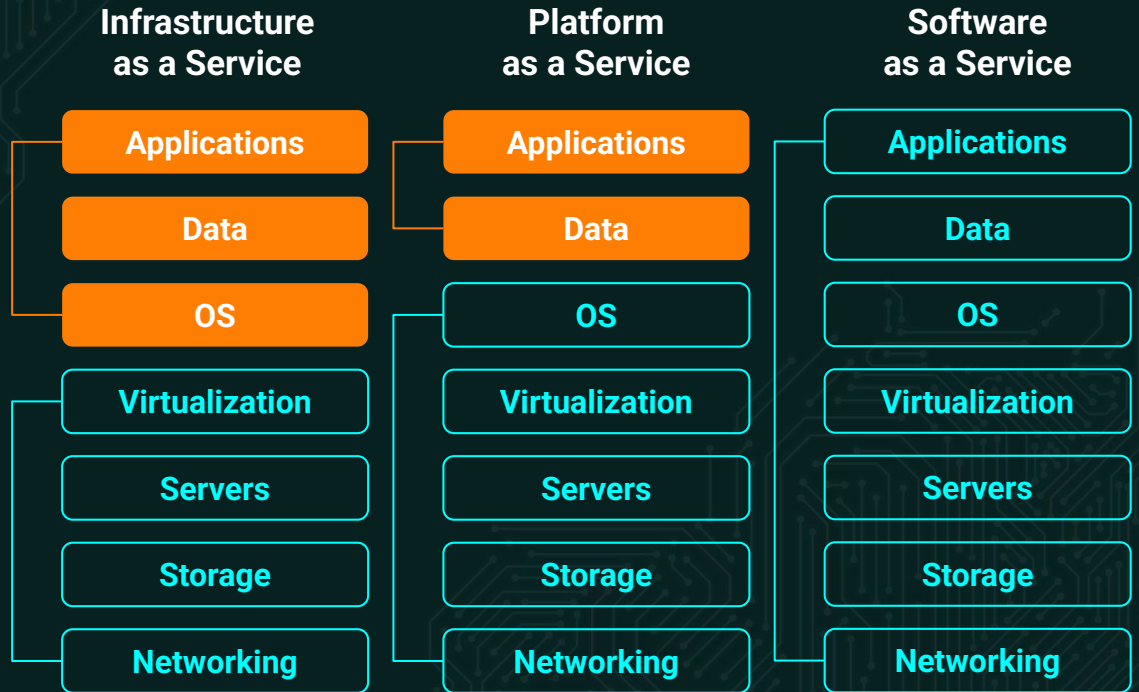
Twenty years later, nothing has changed. At scales on infrastructure. The reason a startup can build with frontier models today, without a data center or a nine-figure budget, is the same reason a startup could launch in 2006 without owning a single server. The foundation always matters. And everything starts from what customers really need.

The question is the same as it was 20 years ago — what problem are you trying to solve that you've been told can't be done? Tell me what's holding you back, and let's figure it out together: nowgobuild@amazon.com

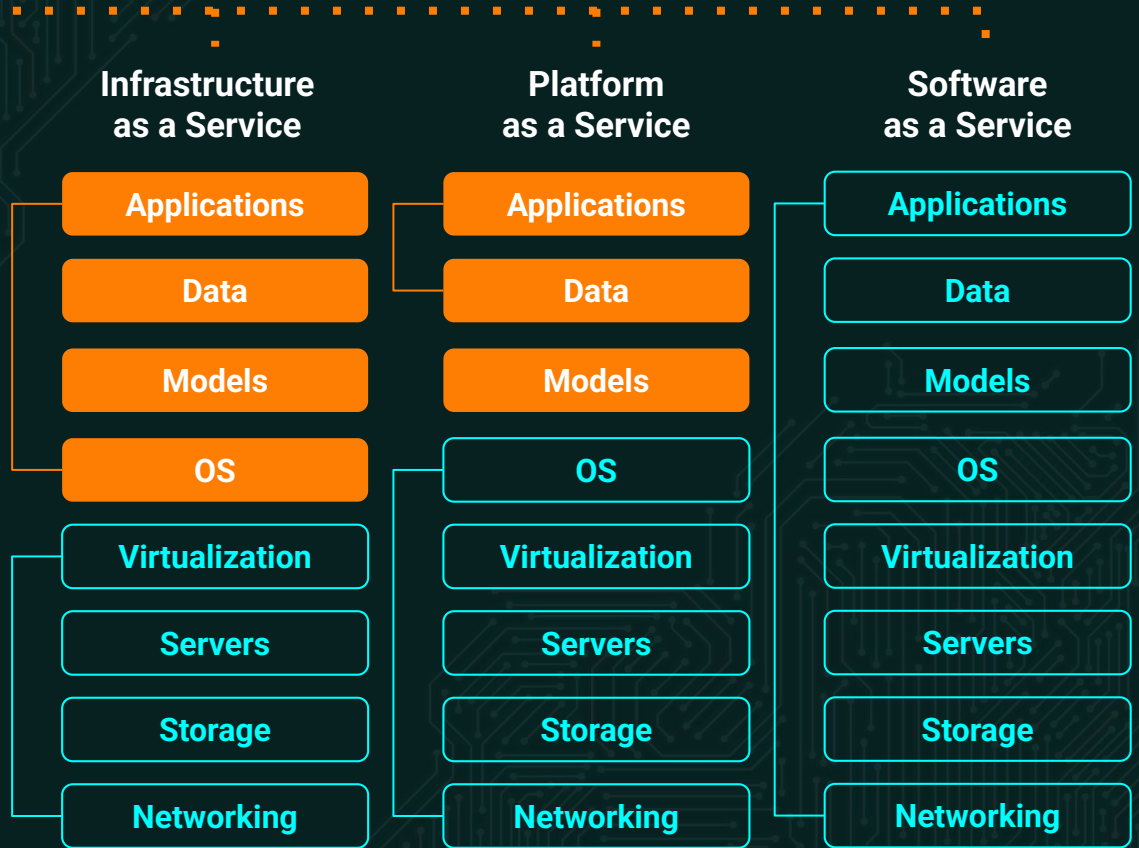
Now, go build.



The shared responsibility model



The shared responsibility model & AI



Private models, SLMs
and retraining

Bedrock, Vercel, Azure
AI Platforms

Chat GPT, Claude &
Gemini Apps



The Current Situation


The same old story, but with some wrinkles.

AI adoption represents a structural continuation of the cloud-native era.

- **Architectural Dependence:** Most AI consumption is primarily cloud-native. The established paradigm of security "in" the cloud remains the foundational rule for AI governance.
- **Defined Accountability:** Organizations are strictly responsible for the layers within their control. Security obligations are dictated by the specific boundaries of your configuration and access.
- **The AI Wrinkles:** AI introduces new challenges to the cloud model. Security teams lack direct visibility into underlying model weights and have limited control over probabilistic outputs.
 - **You don't control the models in most cases**
 - **You don't controls the outputs**



The 5 Layers of AI Security



The Data Layer

Information feeding into or generated by AI. It covers prompts, context data, training sets, and system logs.

The Shadow Layer

Unauthorized AI use within an organization. This includes unvetted SaaS apps, browser tools, and code.

The Identity Layer

Access rights for users and AI agents. It governs permissions to access data and other systems.

The Integration Layer

Systems that connect the AI model to other tools. APIs, autonomous agents, and other applications.

The Model Layer

Core intelligence engines of AI applications. This layer includes foundational model, cloud-hosted APIs and private, self-trained instances.




The Model Layer

Understanding security and data handling of adopted models.

- **Uncertainty:** Risks and data policies vary wildly between different model providers.
- **Discovery:** Apply standardized risk assessments to every adopted model.
- **Action:** Create a whitelist of approved models based on security requirements.

Do we understand the risk profile of every model we use?



The Shadow Layer

Unvetted AI adoption across employees, tools and code.

- **Visibility:** AI adoption is already happening at a massive scale and 90% is invisible to security teams.
- **Discovery:** Build a complete inventory of all AI in use across code, cloud and workforce.
- **Action:** Build a complete inventory of all AI in use across code, cloud and workforce.

Do we know what AI tools are being used and where?




The Integration Layer

AI adoption through applications, agents, and APIs.

- **Attack Surface:** Vulnerable APIs and agents create new entry points for attackers.
- **Monitoring:** Organizations need to understand not just what AI is being used but how and why.
- **Action:** Capture all model integrations with full execution context for all applications & agents.

How are we using AI for every use case?



The Data Layer

Oversight of information flows in prompts and training.

- **Leakage:** Sensitive data flows into prompts, fine-tuning jobs, and agent workflows with little oversight.
- **Observability:** Organizations need to see and understand what data is being shared with AI.
- **Action:** Capture detailed logs with automated detections of PII or sensitive data leakage.

What data are we sharing with AI?

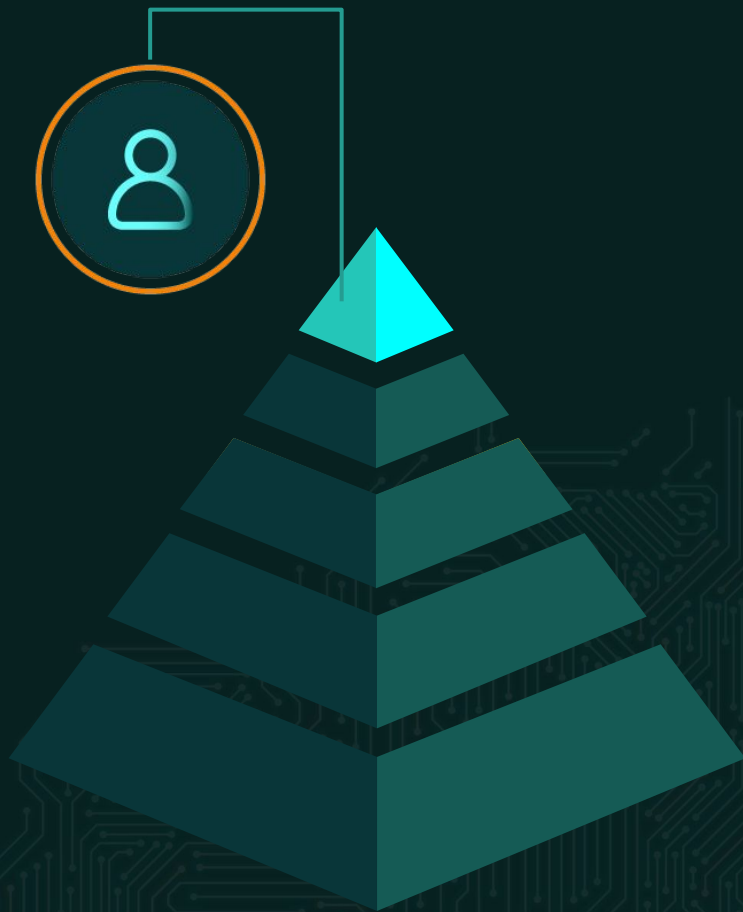


The Identity Layer

Governing permissions for human and agentic entities.

- **Blast Radius:** Over-privileged agents and users turn minor errors into major security failures.
- **IAM:** Manage credentials and roles for non-human AI entities like agents.
- **Action:** Apply the principle of least privilege to every AI agent and user.

**Who is using AI and
what are they using it for?**



The 5 Key Questions at Every Layer



The Identity Layer
Who is using AI and what are they using it for?



The Data Layer
What data are we sharing with AI?



The Integration Layer
How are we using AI for every use case?



The Shadow Layer
Do we know what AI tools are being used and where?



The Model Layer
Do we understand the risk profile of every model we use?

FireTail's Vision .





FireTail helps you say 'yes' to AI.
Identify shadow AI, assess risks,
and build the informed governance
needed to enable AI adoption
without stifling innovation.



Workload AI.


FireTail discovers, assesses and protects AI adoption across development environments. Incorporate AI capabilities into your applications, systems and developer workflows with confidence.



Workforce AI.

FireTail logs, monitors and evaluates AI usage by users across the org. Our managed browser extension helps you to govern and control AI use without stifling innovation or efficiency.

KEY CAPABILITIES



**AI Security.
Solved.**



FireTail Answers All 5 Questions



The Identity Layer
Who is using AI and what are they using it for?



The Data Layer
What data are we sharing with AI?



The Integration Layer
How are we using AI for every use case?



The Shadow Layer
Do we know what AI tools are being used and where?



The Model Layer
Do we understand the risk profile of every model we use?

Jumpstart your AI governance journey.

Get complete visibility in just 15 minutes.

Scan the code and schedule a demo to see how FireTail can help you achieve informed AI governance to help:

- **Get ahead of board risk**
- **Enable experimentation & promote adoption**
- **Minimize risk and protect against breaches**



Thank you.

Come see us at **Booth ESE-52**
www.firetail.ai

