Metropolitan Washington Council of Governments

UTILIZING AI WITHIN LOCAL GOVERNMENT

Metropolitan Washington Council of Governments

AI DAY

AT-A-GLANCE

8:30 AM — Registration & Breakfast

9:00 AM — Welcome & AI Pulse Kickoff

9:15 AM — AI in Cybersecurity – *Palo Alto Networks*

10:15 AM — Break

10:30 AM — AI in Public Safety – *Peregrine*

11:30 AM — Lunch

12:00 PM — Board of Directors Meeting

12:10 PM — Introduction to AI in Local Government

12:15 PM — AI in Government – *Congressman Don Beyer*

1:00 PM — County Compass Toolkit – *NACo*

1:30 PM — AI Policy Development – *HGAC*

2:10 PM — Regional Best Practices Panel

3:15 PM — Bringing It All Together – *Joe Paul & AWS*

4:00 PM — Adjourn

TECHNOLOGY

STARTUPS TO WATCH

These 11 upstarts are transforming their industries and making a mark on D.C.'s innovation scene.

WASHINGTON BUSINESS JOURNAL

OptimaNova AI

About the company: Amid a market flooded with hundreds, if not thousands, of commercially available AI tools, OptimaNova AI LLC hopes to match companies with the ones that are most suitable to maximize the technology. Joe Paul, the former CEO of free computer training and certification provider Byte Back, founded the Northeast D.C.-based company in November 2023 to prepare and train organizations on how AI technology can best be used across a business. OptimaNova determines this by having its customers complete a free assessment test to shed insights on the types of commercial or bespoke AI tools that can best aid a company. The startup then offers consultation and training services on these tools to find ways to implement the technology directly for its customers or partners, which include consumer goods giant Procter & Gamble Co. and D.C. marketing firm The Brand Guild. OptimaNova has 23 employees in full-time and advisory-related roles and has not yet raised any outside funding.

Why we're watching: Over the next year, OptimaNova plans to launch new AI-based products specifically tailored to nonprofits and government organizations. It's also eyeing its first possible outside investment opportunity to help reach its next phase of growth. OptimaNova is on track to finish the year with over $1 million in revenue, which could reach $1.5 million depending on the outcome of some of its government contracting awards. – *Nate Doughty*

Joe Paul is the founder of D.C.-based OptimaNova AI.
JOE PAUL

OptimaNova AI

AI Pulse Survey

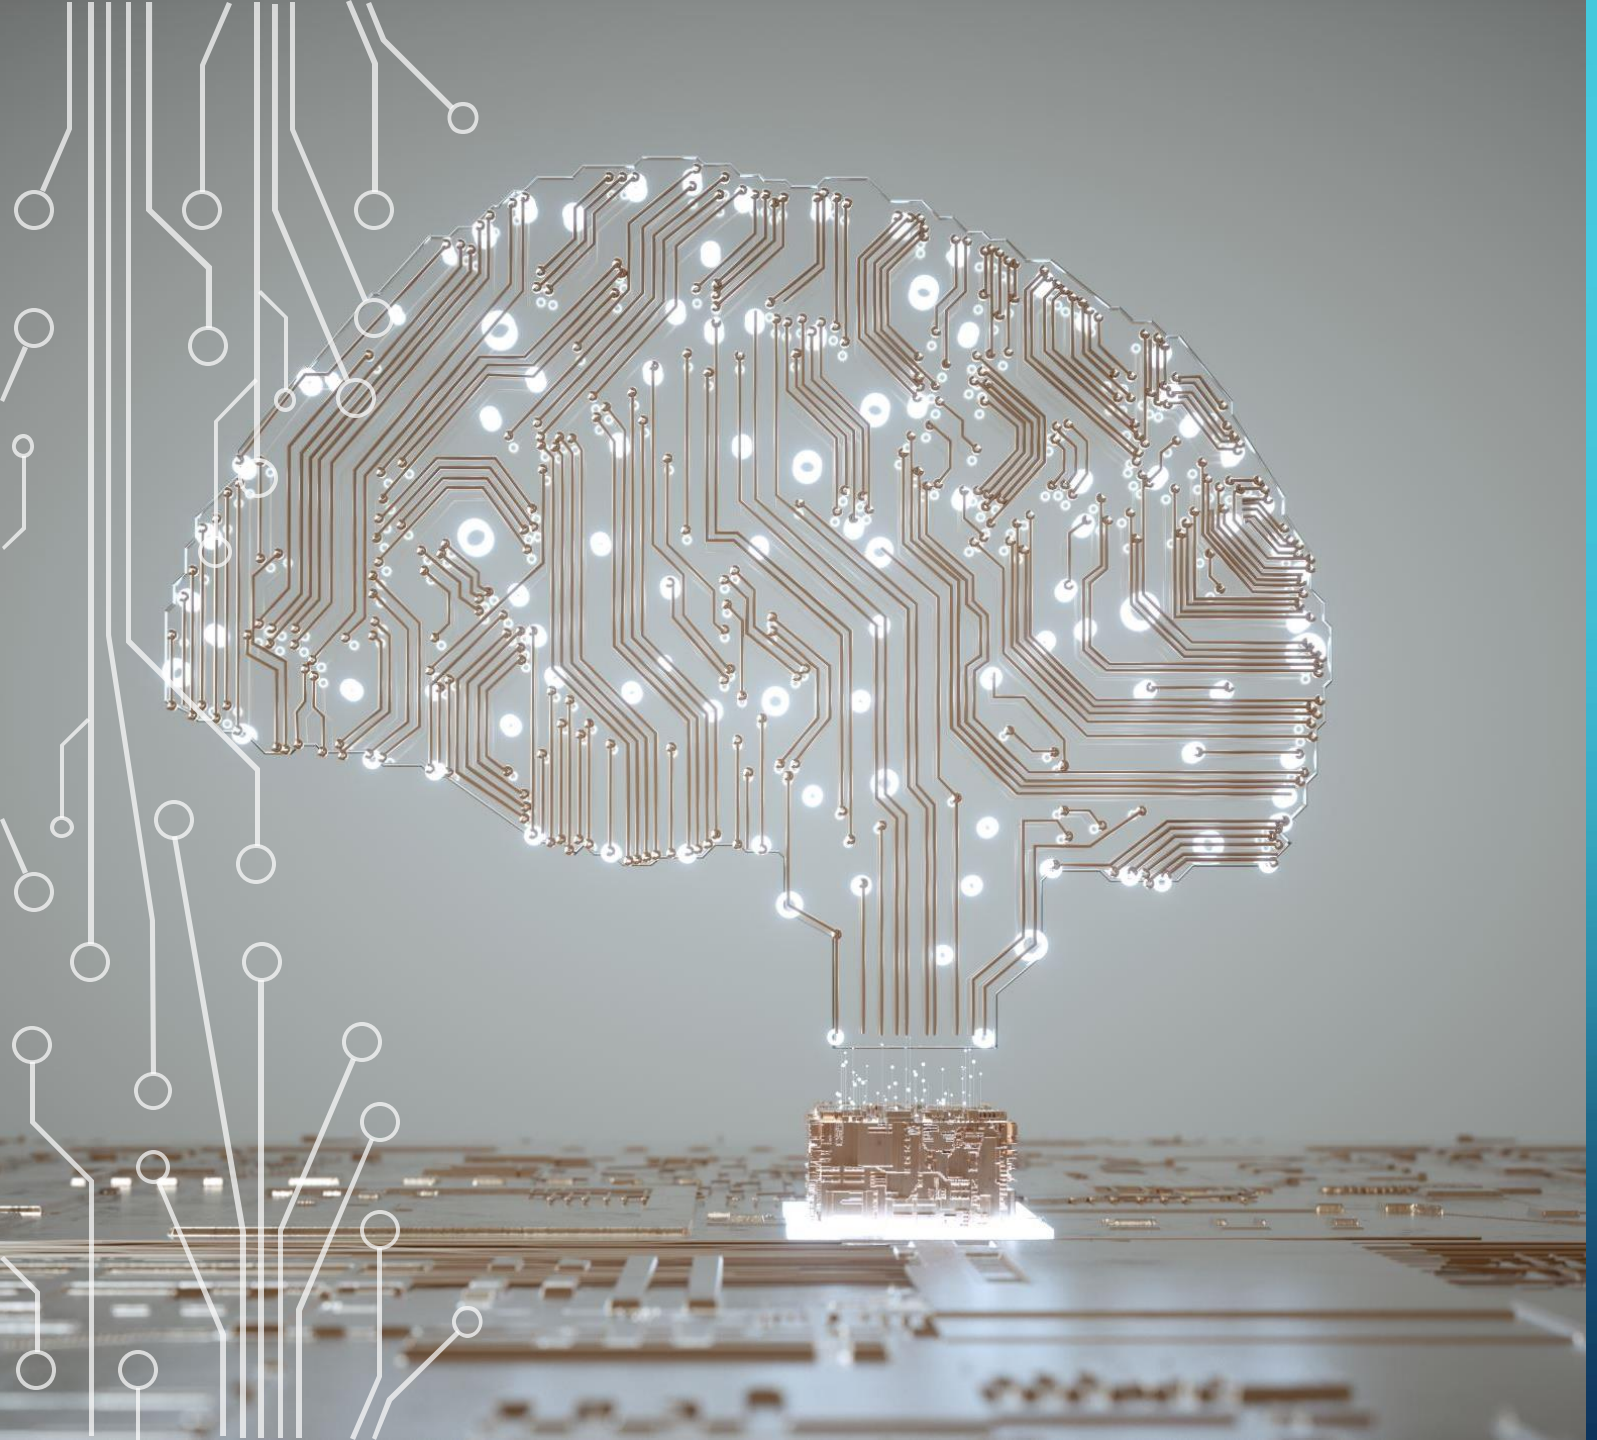

USING AI TO IMPROVE LOCAL GOVERNMENTS

- Explore how AI **strengthens** local government
- From security to **safety**, policy to **practice**, AI is reshaping **public service**
- Our goal: ***Smarter, faster, and more equitable communities***

Metropolitan Washington
Council of Governments

Metropolitan Washington
Council of Governments

Agencies using **AI** in casework could cut
35% of costs within a decade

(Boston Consulting Group 2025)

Only **26% of government organizations** have fully integrated AI—despite **64%** saying it's transformative

(Ernst & Young 2025)

62% cite data security as the *main* barrier

(Ernst & Young 2025)

AI IN CYBERSECURITY

AI security tools detect intrusions
92% faster than legacy systems
(Stanford/Arxiv 2024)

78% of U.S. government IT systems
now use AI in cybersecurity
(SQM 2024)

Metropolitan Washington
Council of Governments

Security for AI?

Jessica Souder, Public Sector Lead, Prisma AIRS

Version 1.1

September 2025

First drug created by AI enters clinical trials

GlobalData Healthcare

Unlike other AI-produced drugs in trials, INS018_055 is the first drug with both a novel AI-discovered target and a novel AI-generated design.

Recommended Reports

Artificial Intelligence in Pharmaceuticals: In-silico drug d...

Innovation in Pharmaceuticals: Transcription factors for AAV

Innovation in Pharmaceuticals: Embryonic stem cell culturing

LOA and PTSR Model - Drugs for Neurodegenerative Diseases

Artificial Intelligence (AI) in Drug Discovery - Thematic Re...

View All

SCIENCE / TECH / ARTIFICIAL INTELLIGENCE

AI suggested 40,000 new possible chemical weapons in just six hours

An instructor at the Fort Leonard Wood Chemical School, who is designated as an agent handler, carries the VX nerve agent to contaminate a jeep in one of the eight chambers used for training chemical defense on April 18, 2003 at Fort Leonard Wood, Missouri. Photo by Brendan Smialowski/Getty Images

/ 'For me, the concern was just how easy it was to do'

By Justine Calma, a science reporter covering the environment, climate, and energy with a decade of experience. She is also the host of the Hell or High Water podcast.

Mar 17, 2022, 1:06 PM MDT | 0 Comments / 0 New

New Zealand

Supermarket AI meal planner app suggests recipe that would create chlorine gas

Pak 'n' Save's Savey Meal-bot cheerfully created unappealing recipes when customers experimented with non-grocery household items

Tess McClure *in Auckland*

@tessairini

Thu 10 Aug 2023 00.19 EDT

What if your beer started producing cyanide?

Baselining our collective knowledge

- Who feels comfortable with the term model?
- What is an application versus a model?
- What are third-party repositories? Have you heard of Hugging Face?

Do NOT feel bad if you don't….let me tell you about a talk I had with an FBI SES…

The Answers

What is a *Model*?

- A **model** is the *core mathematical component* of an AI system.
- It's trained on data, contains parameters/weights, and performs inference (mapping inputs → outputs).

Examples:

- GPT-4 (an LLM model)
- ResNet (an image classification model)
- XGBoost (a gradient boosting model for tabular data)

Key idea: Models don't "stand alone" in production — they're usually embedded inside something larger.

What is an *Application*?

- An **application** is the *system that uses one or more models to deliver functionality* to end users.
- It includes:
  - The **model(s)**
  - APIs, user interfaces, and business logic
  - Supporting infrastructure (databases, logging, access controls)
  - Integration into other enterprise services

Example:

- A fraud-detection **model** (logistic regression trained on transactions)
- Inside a fraud-detection **application** (that connects to banking apps, flags suspicious activity, alerts analysts).

To explain it a bit more…

Analogy

Think of it like **engines vs. cars**:

- The *engine* = the **model** (power source, technical core).
- The *car* = the **application** (engine + chassis + wheels + user controls).
- You don't drive an engine by itself — but the engine's reliability is critical to the car's performance.

Bottom Line

- *Models* are the mathematical brains.
- *Applications* are the systems built around those brains.
- Security folks often think "application-first," while ML folks think "model-first." That's why framing matters — using the right term depends on the audience.

What are third-party repositories - like Hugging Face?

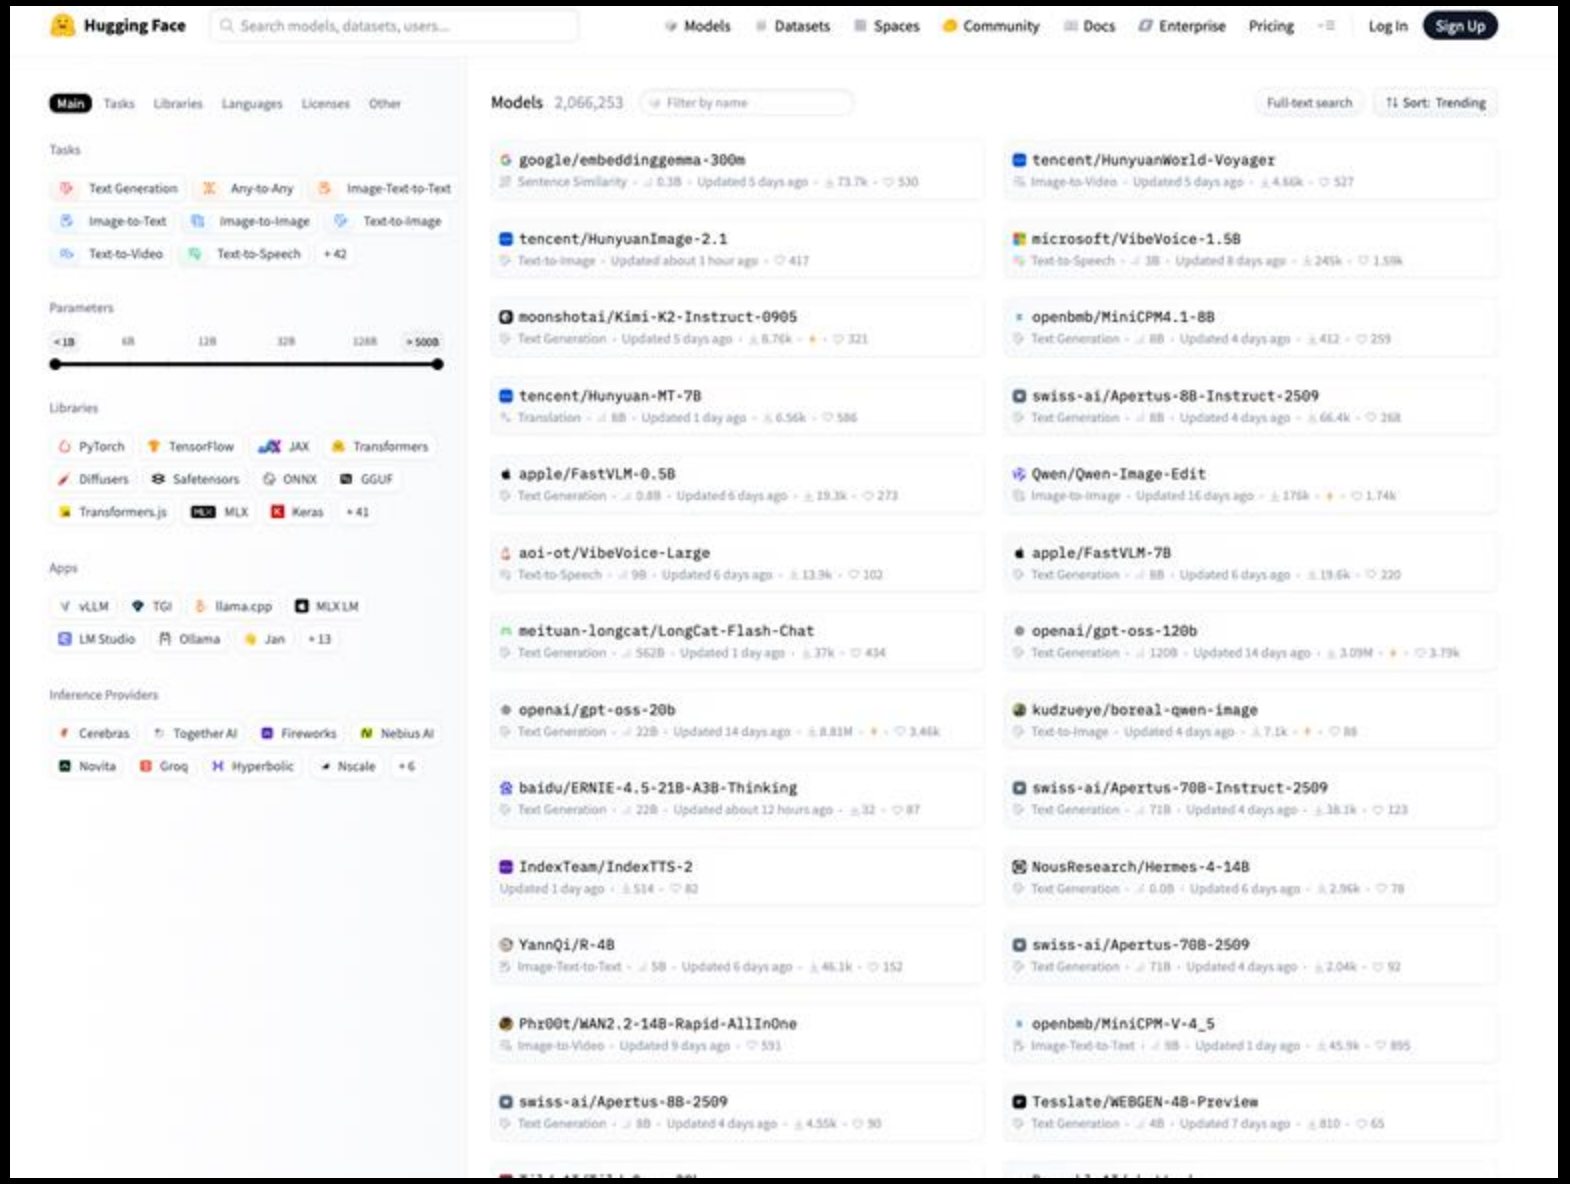

Models are essential for building AI apps…

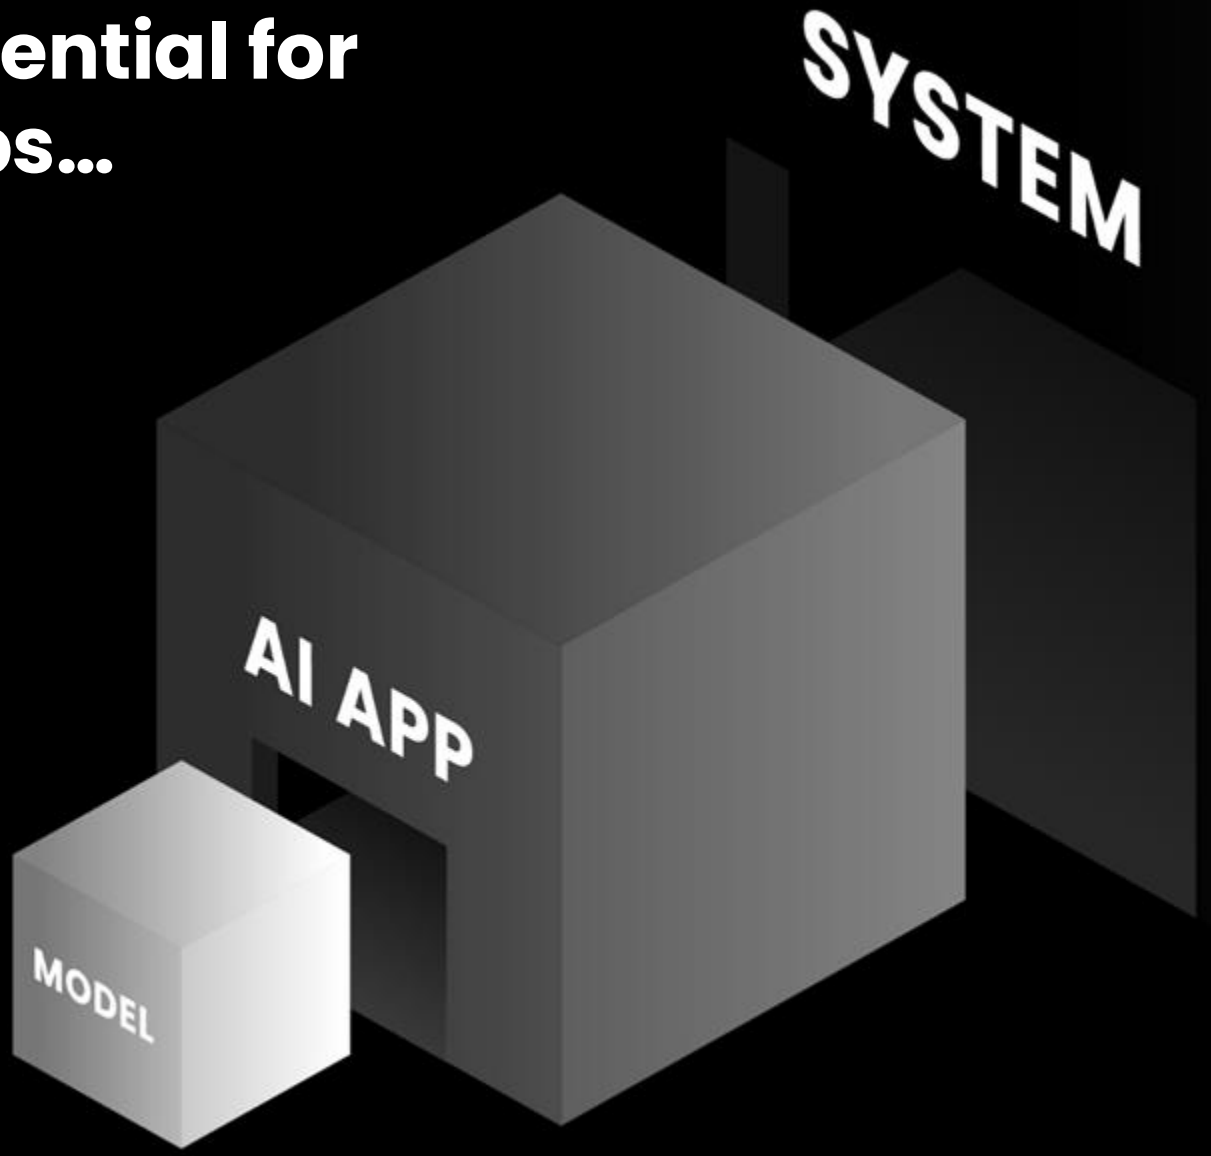

But… hidden threats leave your systems vulnerable.

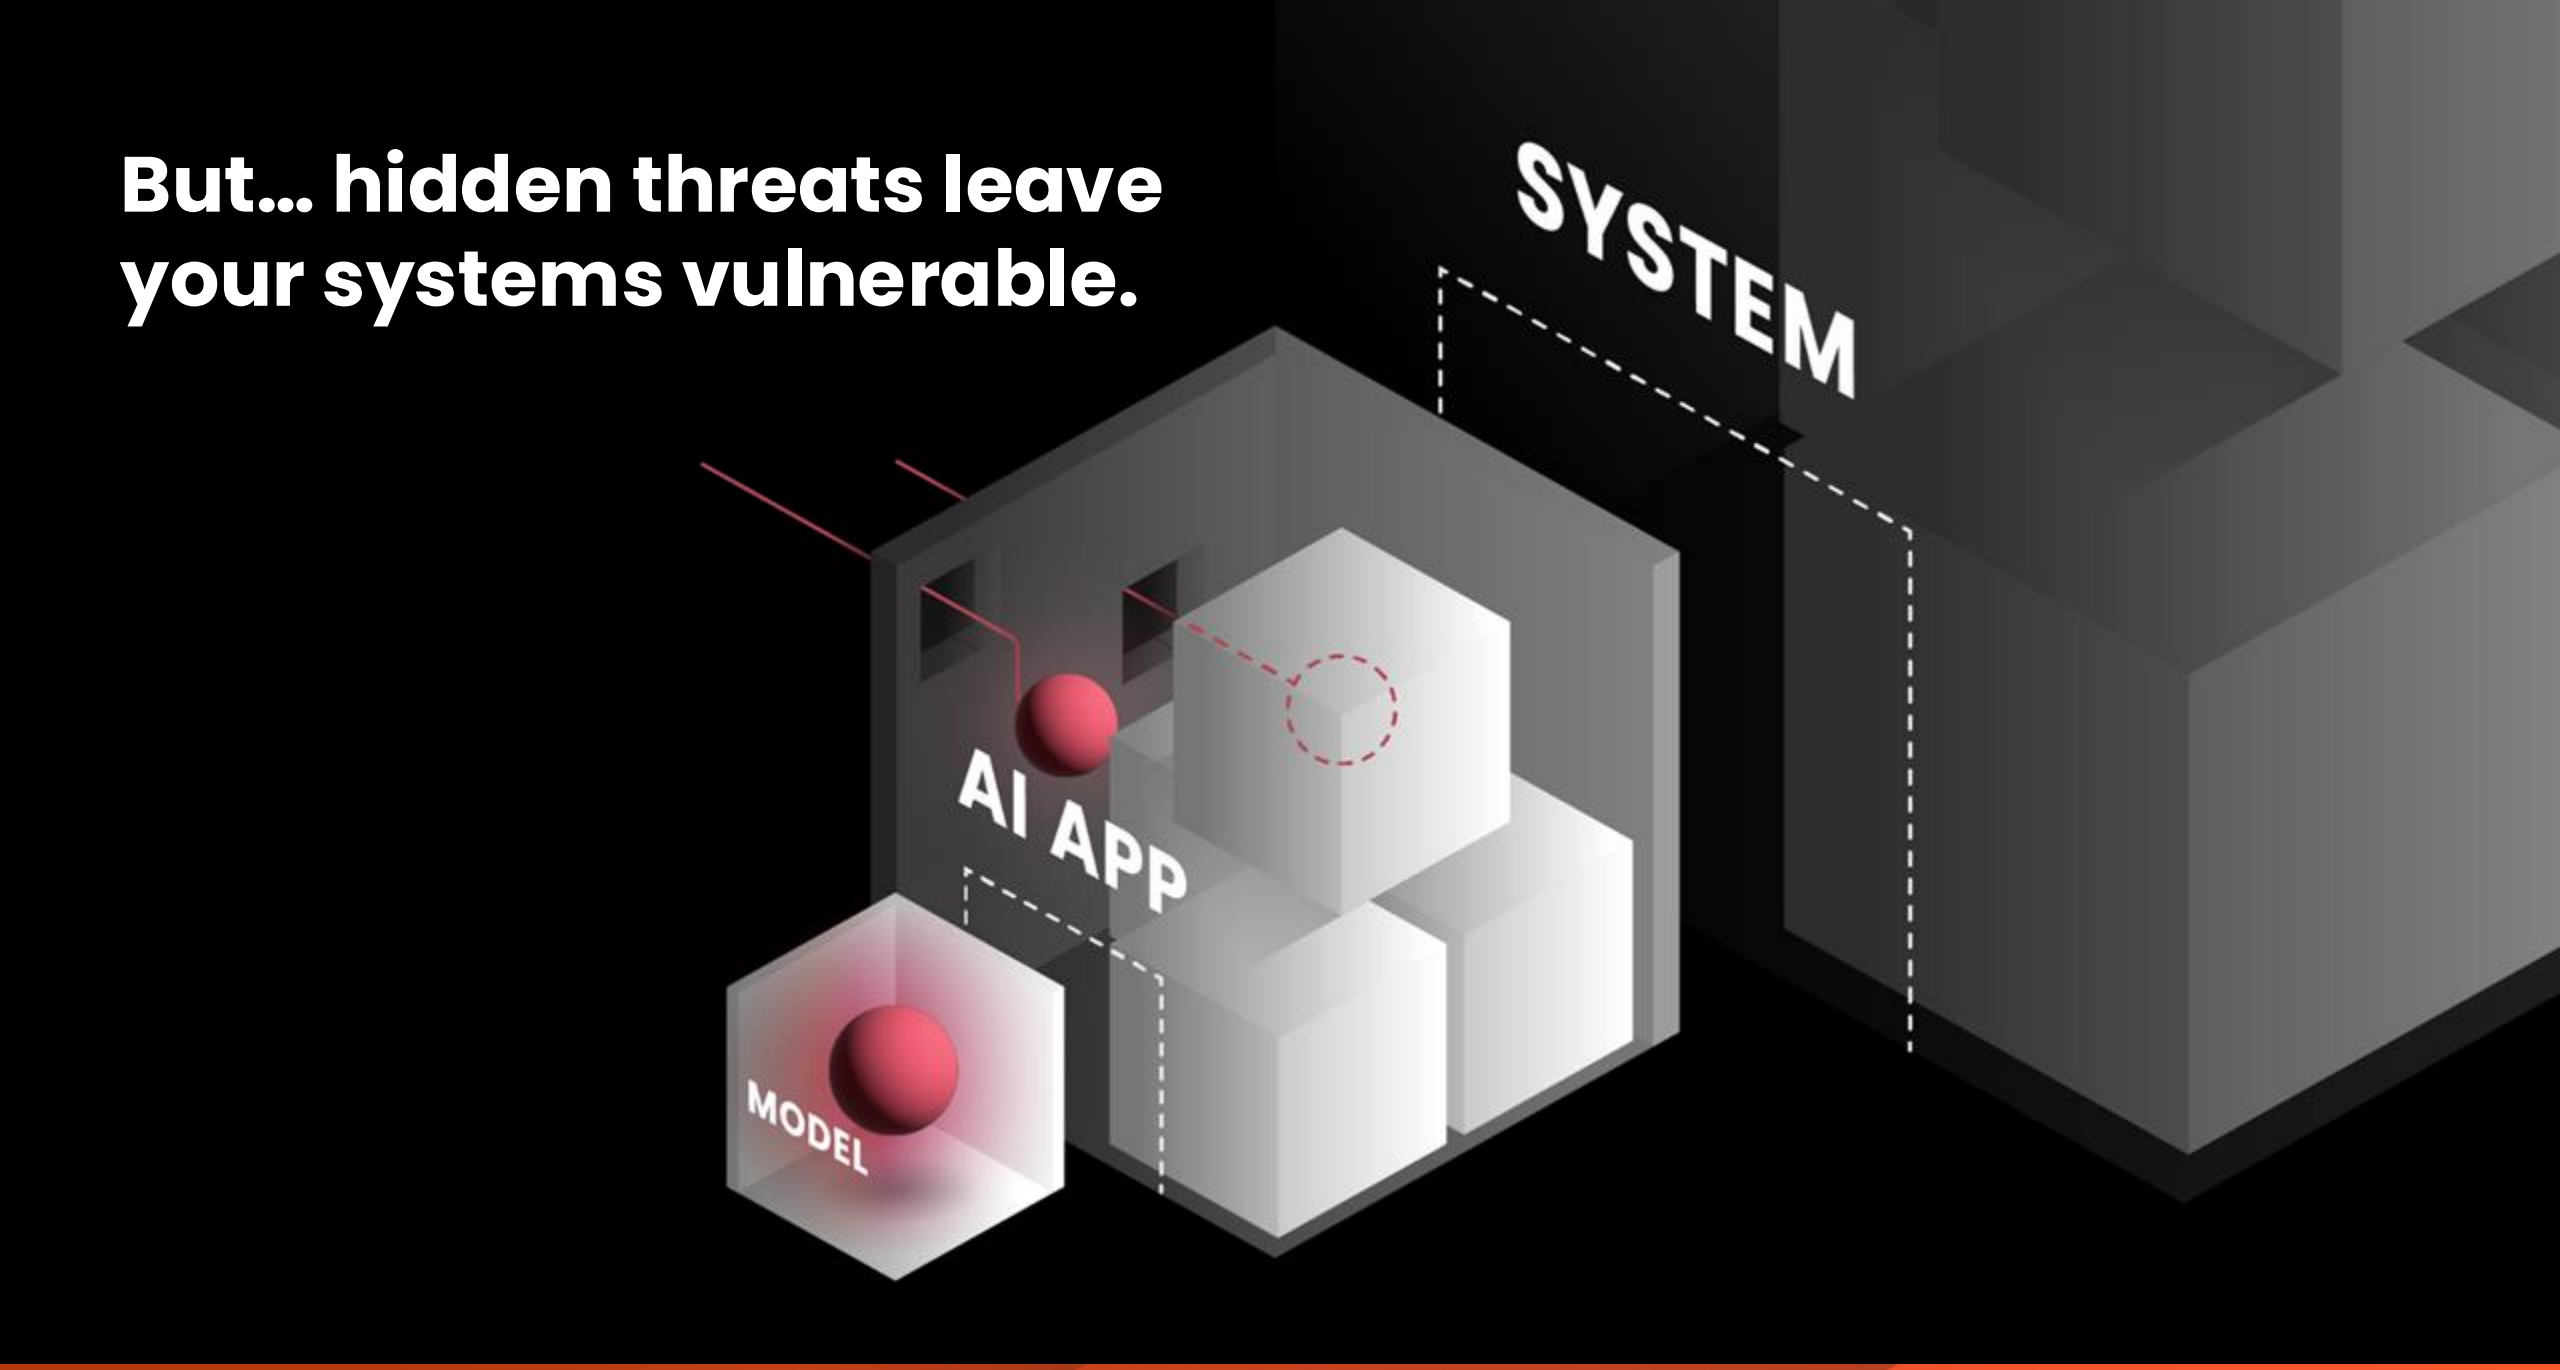

How many models do you think are in production *right now*?

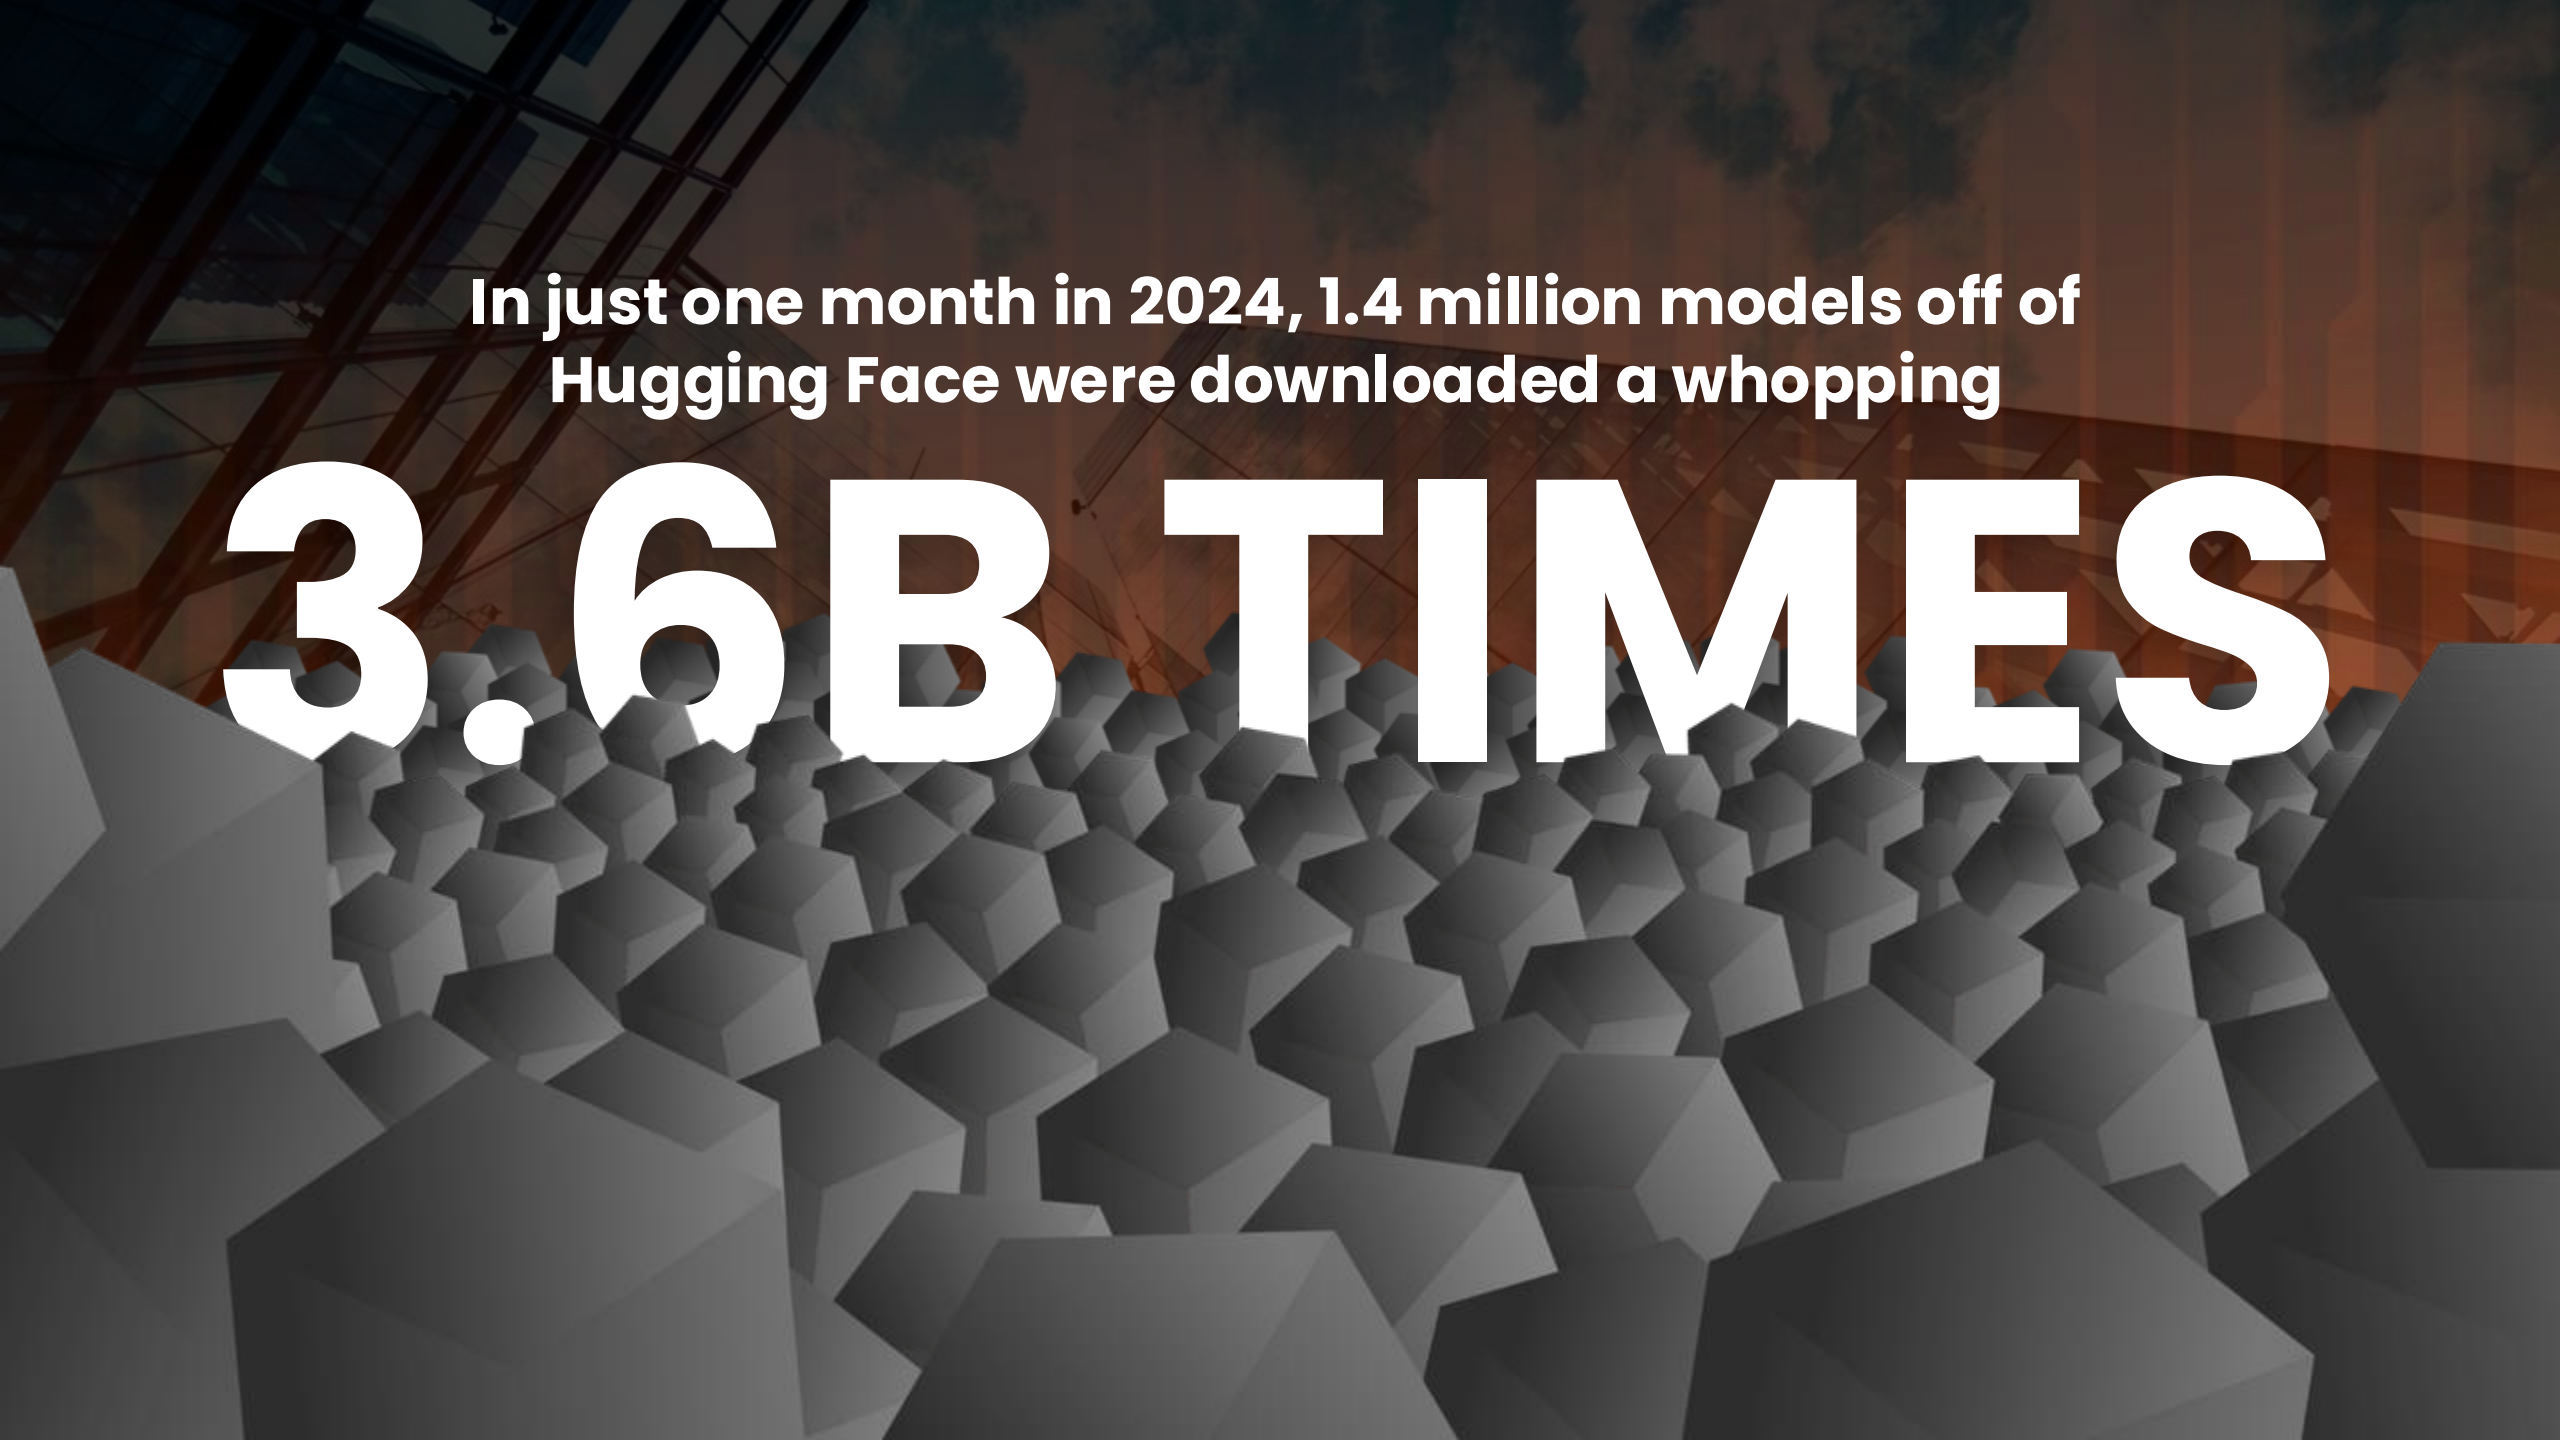
In just one month in 2024, 1.4 million models off of Hugging Face were downloaded a whopping
3.6B TIMES

Model Files

- Most model files can execute code
- Most antivirus do not detect malicious model files
- Multiple paths of attack:
  - Attached to phishing email
  - Uploaded to model repositories like HuggingFace
  - Uploaded to vulnerable MLOps tools
    - Hundreds of MLOps tooling vulnerabilities have been found by us in the past 2 years

Model Security

Model Files are Invisible Viruses

Dan McInerney January 24, 2024 • 4 minute read

The Underestimated Risk of Model Files in Machine Learning

When a Machine Learning (ML) model is trained it is stored in memory. To save it to disk, so it can be shared with others requires storing it in various formats. The most common and prominent formats, such as pickle, are vulnerable to deserialization attacks where code can be injected into the model which will run upon the model being loaded. This injected code does not affect the model's ability to perform inference, making it difficult to detect malicious models unless specific tools such as Protect AI's Guardian are used. Today's antiviruses and email filters don't detect payloaded model files making these the perfect phishing campaign attachment. Move over PDFs and macro-enabled Word documents, model files are the new kingphisher.

Predictive AI Threat Surface

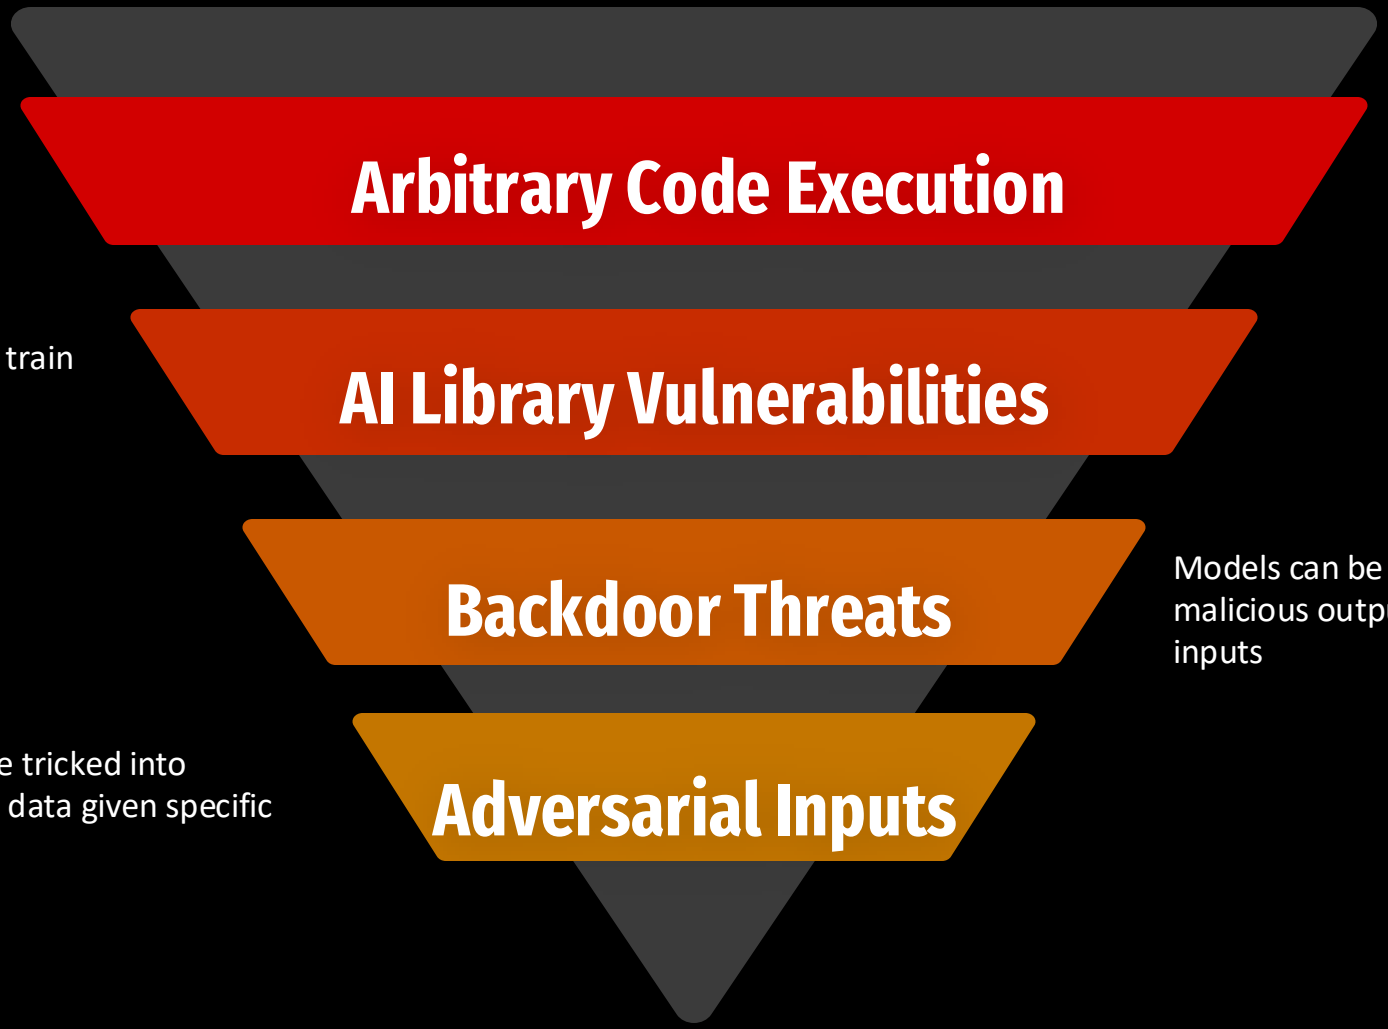

Model files can execute arbitrary code upon being loaded

Insecurity in libraries used to train and track models

Models can be payloaded to trigger malicious outputs given specific inputs

Models can be tricked into misclassifying data given specific inputs

Generative AI Threat Surface

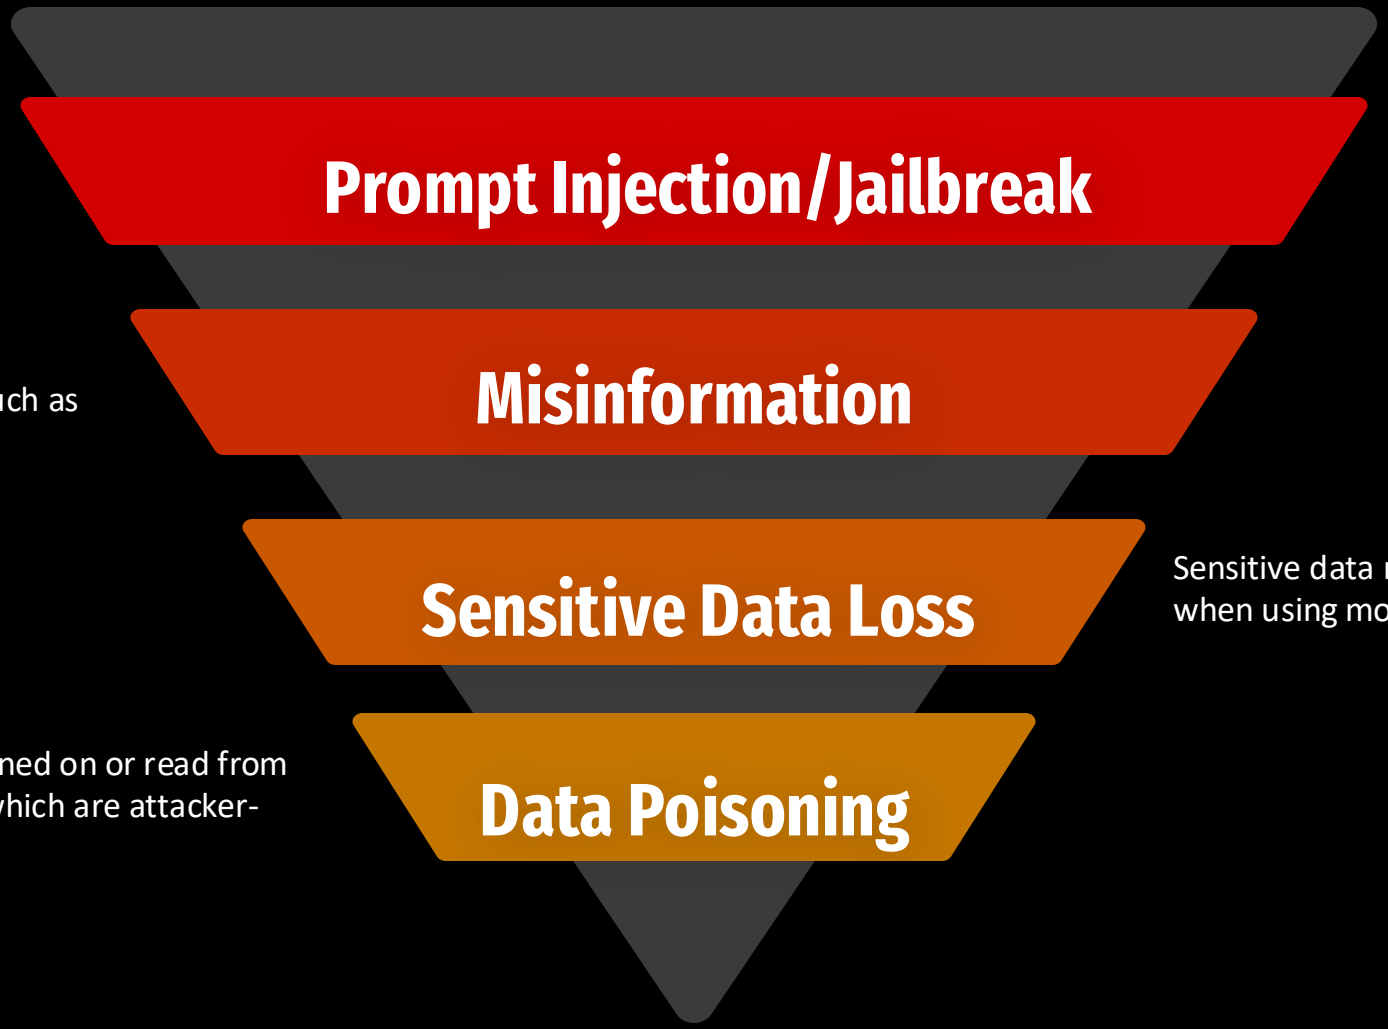

LLM applications can be tricked into bypassing safeguards or returning attacker-controlled output

LLMs may hallucinate facts; especially impactful in fields such as medicine, law, and science

Sensitive data may be sent to 3rd parties when using model providers' API

LLMs may be trained on or read from sources of data which are attacker-controlled

Two Ways Enterprises are Using AI Today

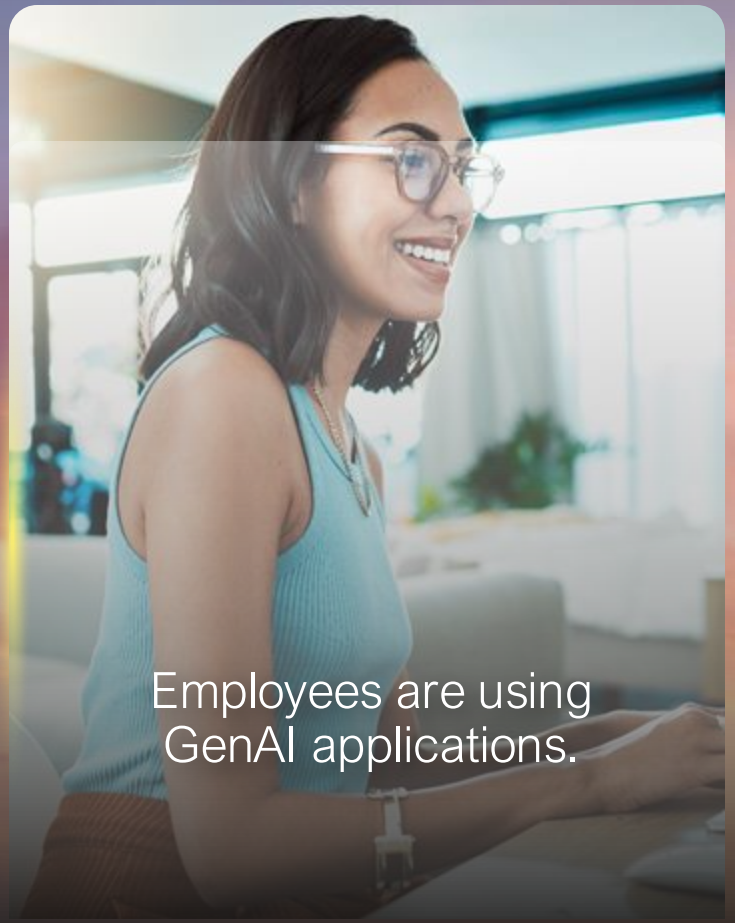

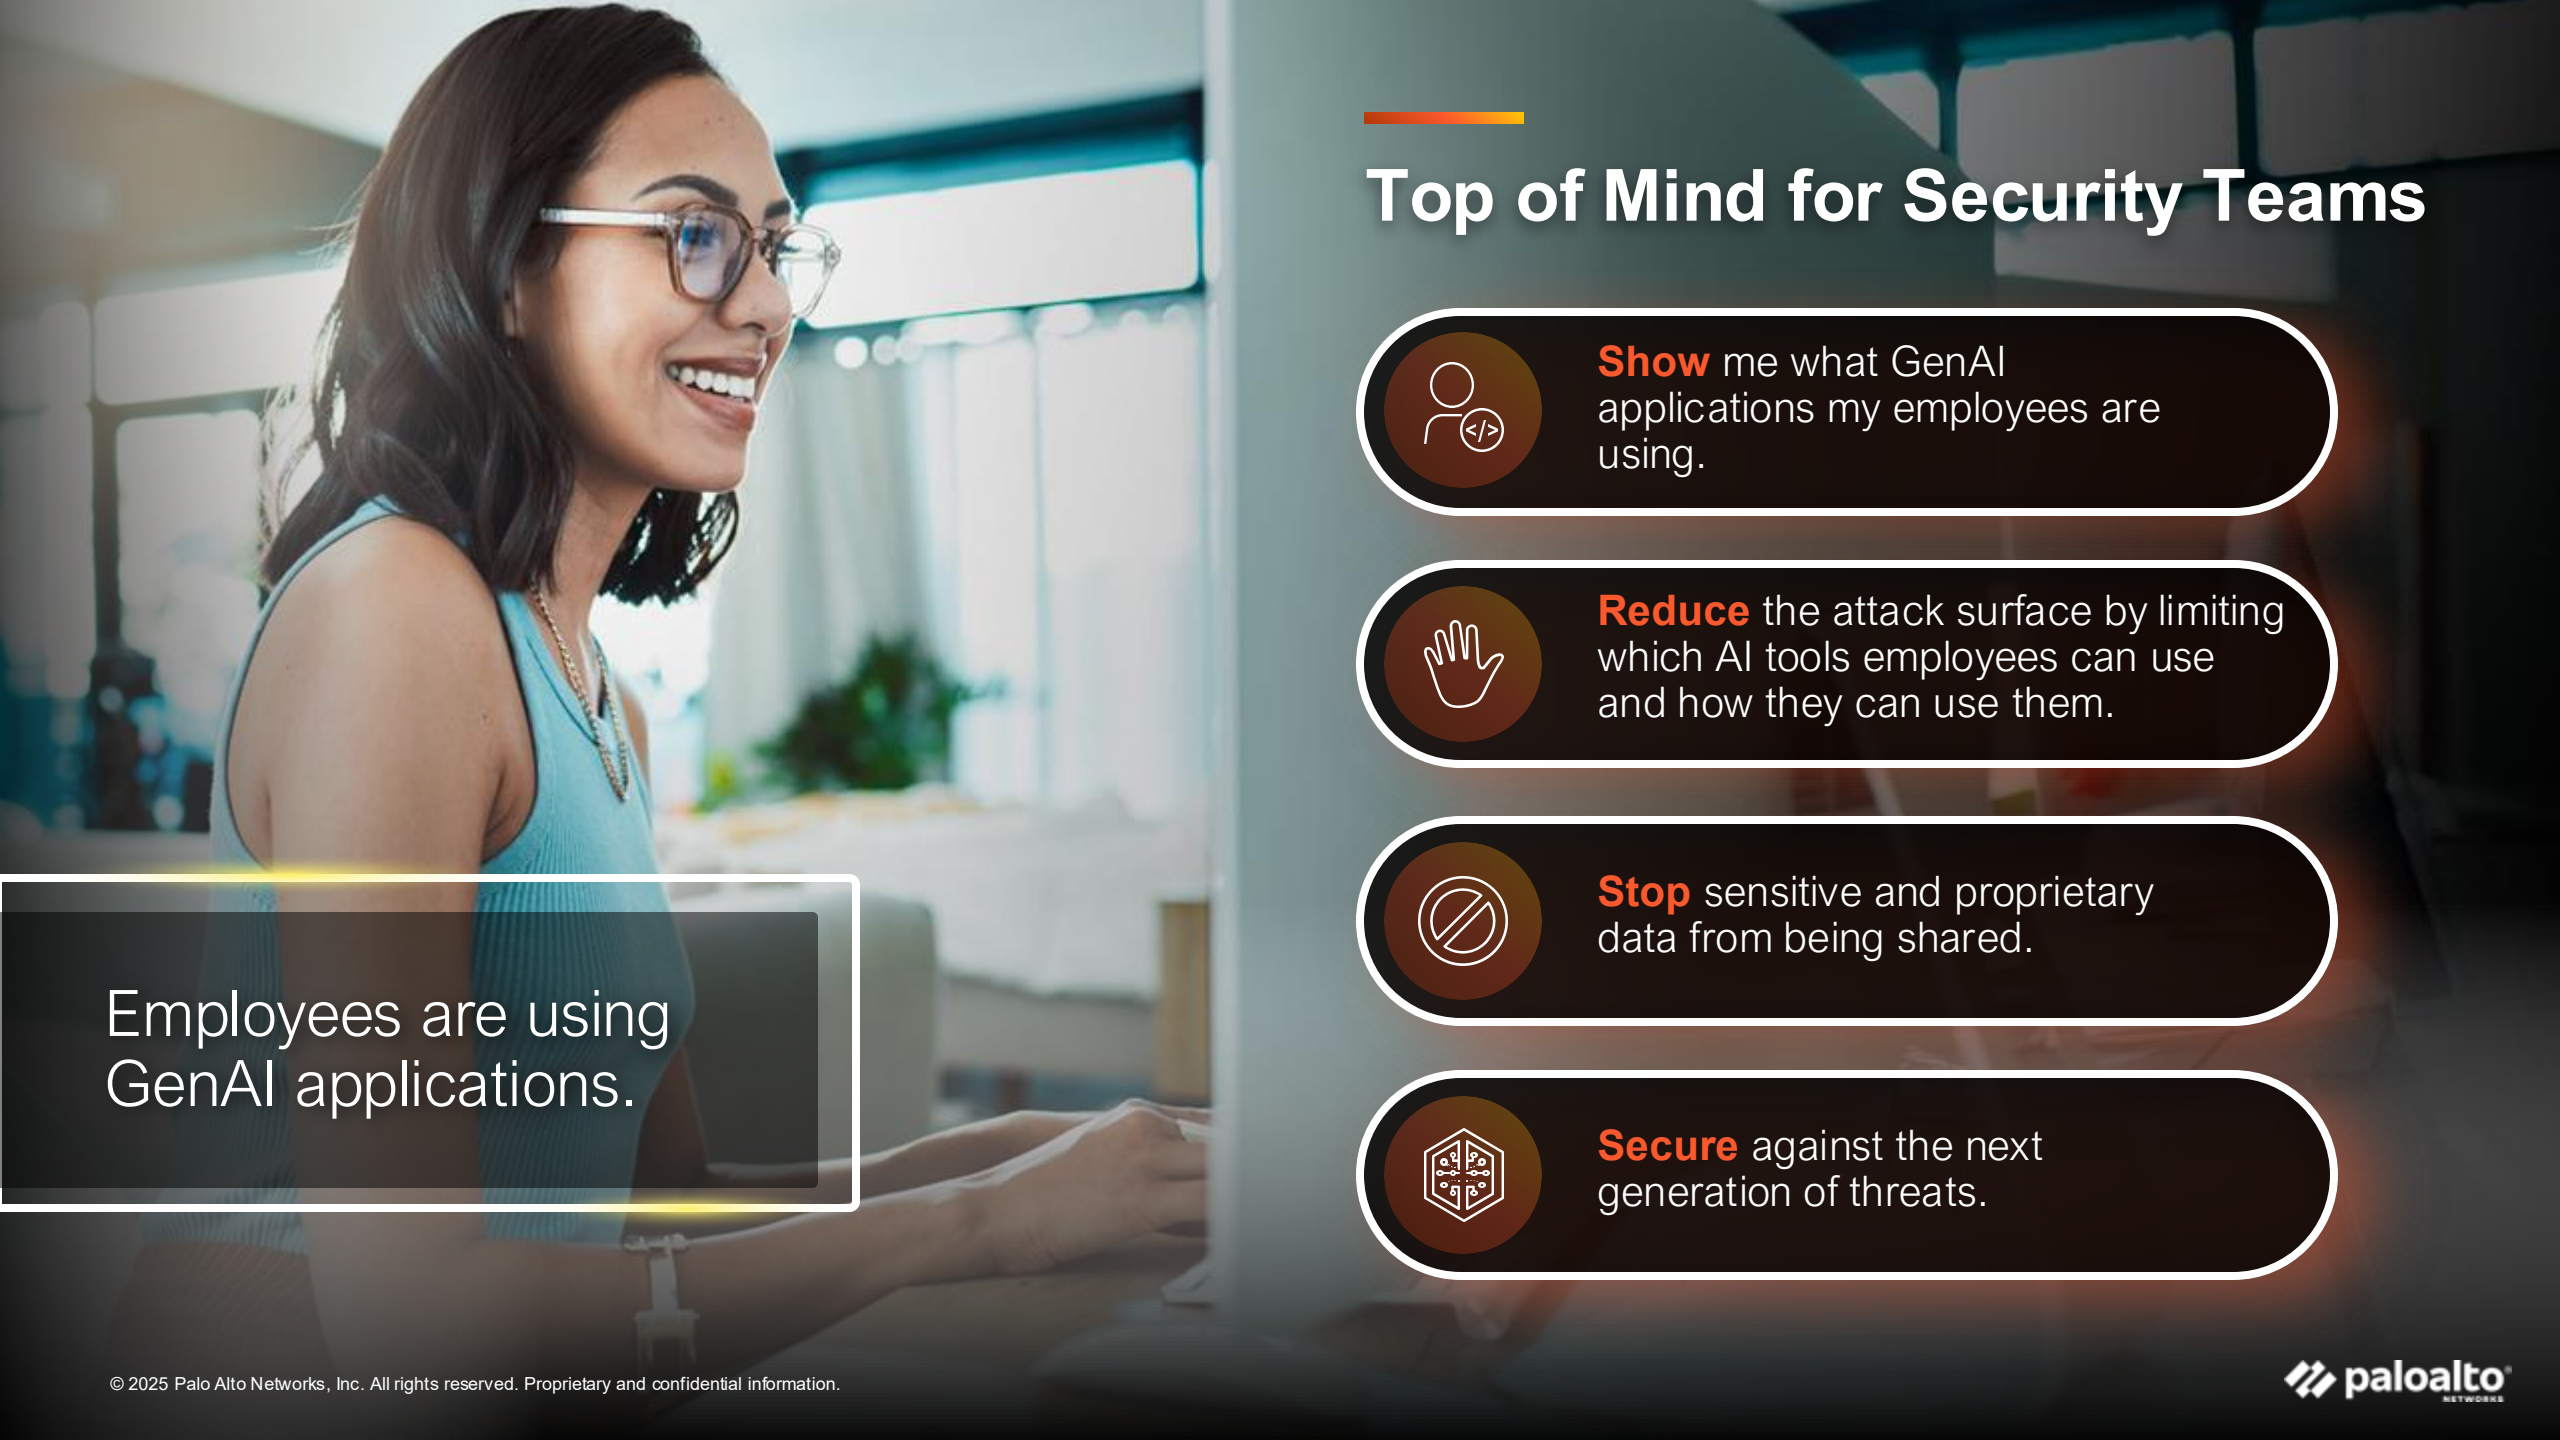

Employees are using GenAI applications.

Top of Mind for Security Teams

- **Show** me what GenAI applications my employees are using.
- **Reduce** the attack surface by limiting which AI tools employees can use and how they can use them.
- **Stop** sensitive and proprietary data from being shared.
- **Secure** against the next generation of threats.

paloalto NETWORKS

AI Apps Bring New Risks

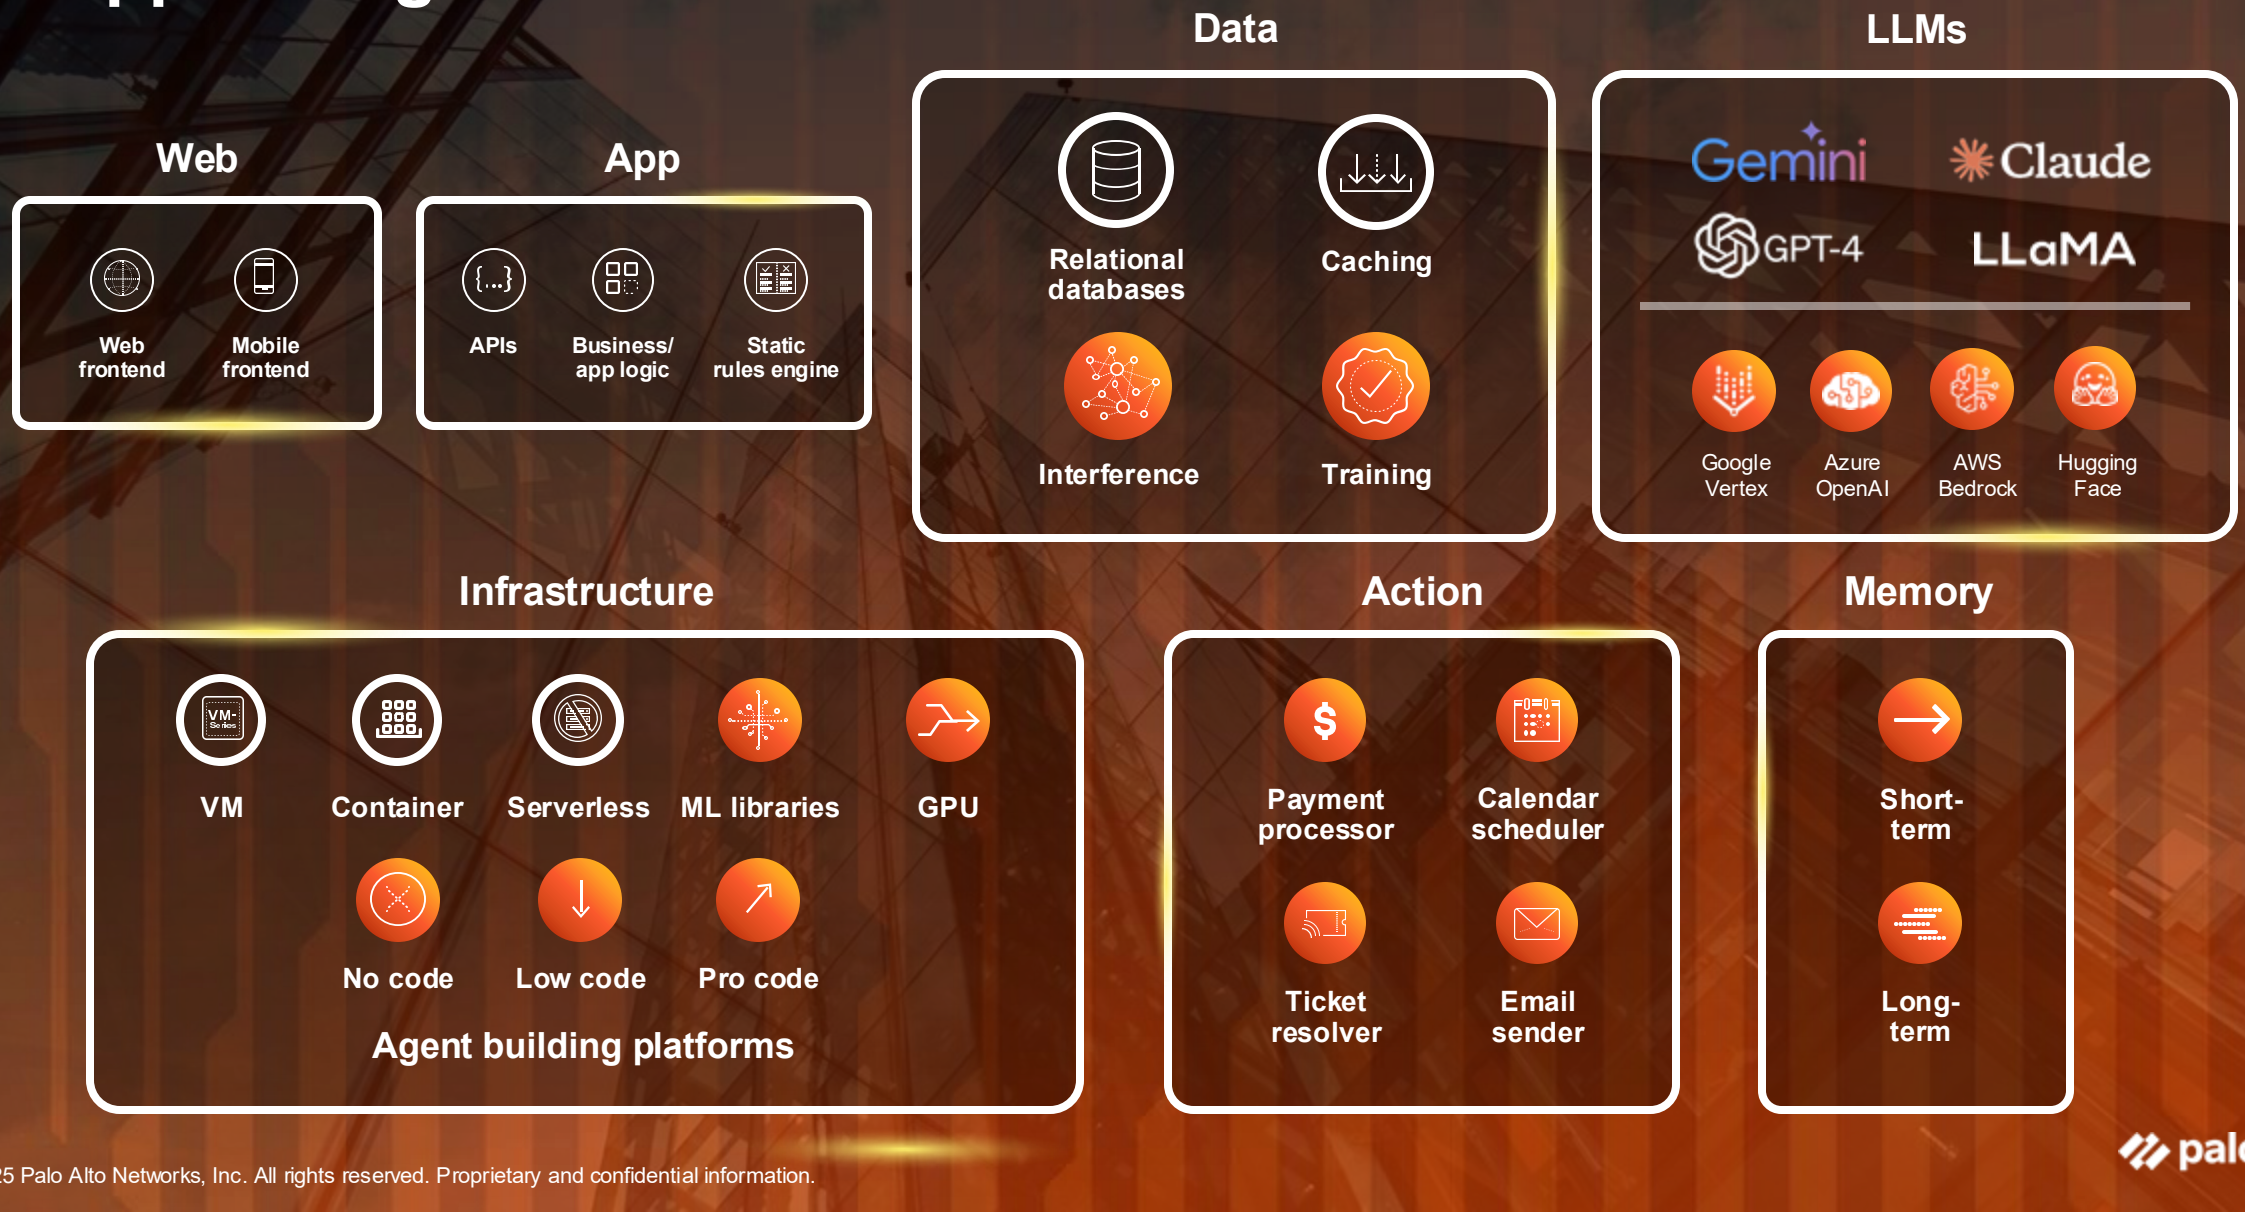

…Adding New Risks…

Web

- Access
- Authentication

App

- Malware Execution
- Unsafe URL Processing

Data

- Publicly Writable Dataset
- Sensitive Data Exposure
- Unrestricted DB Queries

LLMs

- Deserialization
- Insecure Output
- Prompt Injection
- Model DoS Attacks

Infrastructure

- Lack of Segmentation
- Access Control
- Platform Misconfiguration

Action

- Tool Misuse
- Goal Manipulation
- Identity Impersonation
- Excessive Permissions

Memory

- Memory Poisoning
- Privilege Compromise

This is what can be done to secure AI.

AI Risks & Concerns:

Supply Chain
Can I track my AI artifacts?
What do my AI assets contain?

Privacy
What data might be exposed?
What data must never be shared?

Forensics
What happens to the right of "Boom"?
Where are the logs on my assets?

Threat to Remediation Cycle
How do I stay informed on the threats?
What can I do to stay secure & safe?

Environment Visibility
What are people using?
How are they using it?

Risk Impact & Prioritization
Which risks are most important?
What remediations do I take first?

We are *the enterprise standard*
to see, know, & *manage AI risks.*

The result: Secure the most sensitive & critical AI

Finance

Healthcare

Technology

Government

Strengthening AI vuln management for one of the world's largest credit card companies.

The Need
Enhance VLM as they handle massive amounts of sensitive customer data via their extensive credit card operations and global customer services.

Our Solution
A stringent VLM-focused security posture built using:

- Integrated model scanning before deployment
- Automated threat remediation
- Automated red teaming tailored to GenAI systems
- Real-time visibility that produces actionable insights

Meeting AI compliance standards for one of the world's largest medical testing companies.

The Need
Keep the GenAI applications that their clinicians use across 2,000+ patient centers secure and compliant with strict global and FDA standards.

Our Solution
End-to-end security that enabled ongoing advancements without compromising compliance with:

- Real-time observability of LLMs
- Automated red teaming tailored to GenAI systems
- Rapid response tools to address issues before escalation
- Continuous feedback loops for refinement

Tightening AI security for one of the world's biggest online marketplaces.

The Need

Protect over 150 million active users by enhancing app security for their popular AI-driven platform powered by over 100,000 ML models.

Our Solution

A comprehensive and proactive security framework that focused on:

- Continuous, automated monitoring
- Scrutinizing models before deployment
- Community-sourced vulnerability detection
- Ongoing feedback loops for refinement

Deploying AI offensive security for one of the largest data center & colocation companies .

The Need
Secure the GenAI applications behind their >200 data centers in a way that neutralizes threats *before* they turn into breaches.

Our Solution
An offensive SecOps approach that included:

- Automated red teaming tailored to GenAI systems
- Community-driven threat intelligence
- Real-time visibility that produces actionable insights

THANK YOU!

Attack Surface

Prompt Injection Attacks
Attackers craft inputs that "inject" malicious instructions into the prompt, manipulating the model's behavior or bypassing safety filters.

Jailbreak Attacks
A subset of prompt injection, these are designed to force the model to ignore its built-in ethical or safety guidelines and produce prohibited outputs.

Adversarial Examples
Slightly perturbed or carefully engineered inputs cause the model to generate incorrect, harmful, or unintended outputs.

Model Extraction Attacks
By querying the model extensively (often via public APIs), adversaries attempt to reconstruct a surrogate model or infer proprietary parameters and architecture details.

Membership Inference Attacks
Attackers analyze outputs to determine whether specific data points were included in the model's training dataset, potentially compromising privacy.

Model Inversion Attacks
These attacks aim to reconstruct or reveal sensitive aspects of the training data by "inverting" the model's outputs.

Data Poisoning Attacks
Malicious data is introduced into the training process so that the model learns incorrect or harmful behaviors—this can include backdoor or Trojan triggers.

Backdoor/Trojan Attacks
Similar to data poisoning, but with a focus on embedding hidden triggers that, when activated by specific inputs, cause the model to behave in a controlled (and usually harmful) way.

Evasion Attacks
Inputs are crafted specifically to bypass moderation filters or detection mechanisms, often allowing harmful content to be generated or disseminated.

Adversarial Reprogramming
An adversary repurposes the model to perform tasks it wasn't intended for by carefully designing the input, essentially "reprogramming" the model on the fly.

Watermark Removal or Circumvention Attacks
Techniques aimed at removing or bypassing embedded watermarks or other intellectual property protections that help identify or secure the model's outputs.

AI

IN PUBLIC SAFETY

EMERGENCY RESPONSE TIMES SHRINK 20–35% WITH AI-ENABLED DISPATCH

(AMBIQ 2024)


LUNCH

METROPOLITAN WASHINGTON COUNCIL OF GOVERNMENT

BOARD OF DIRECTORS

Rodney Lusk, *COG Board Chair*
Clark Mercer, *COG Executive Director*

Metropolitan Washington
Council of Governments

AI

IN LOCAL GOVERNMENT

AI Transforming Local Government

Public Service Automation

AI automates routine government tasks, improving service speed and reducing operational costs for local administrations.

Data-Driven Decision Making

AI analyzes large volumes of data to inform better decisions and optimize resource distribution within local governments.

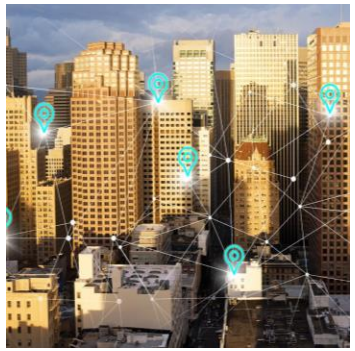

Enhanced Citizen Engagement

AI-powered platforms improve citizen interaction with local government through chatbots, predictive analytics, and smart city projects.

Congressman
Don Beyer

U.S. Representative for Virginia's 8th District

Former Lieutenant Governor of Virginia and U.S. Ambassador to Switzerland and Liechtenstein

Vice Chair of the Congressional AI Caucus; member of the House AI Task Force

Author of AI transparency and research-access legislation; studying machine learning at George Mason University

Metropolitan Washington
Council of Governments

Metropolitan Washington Council of Governments

UTILIZING ARTIFICIAL INTELLIGENCE WITHIN LOCAL GOVERNMENT

National Association of Counties

County Policy Priorities on Artificial Intelligence

Washington Metropolitan Council of Governments Board Convening on AI | September 2025

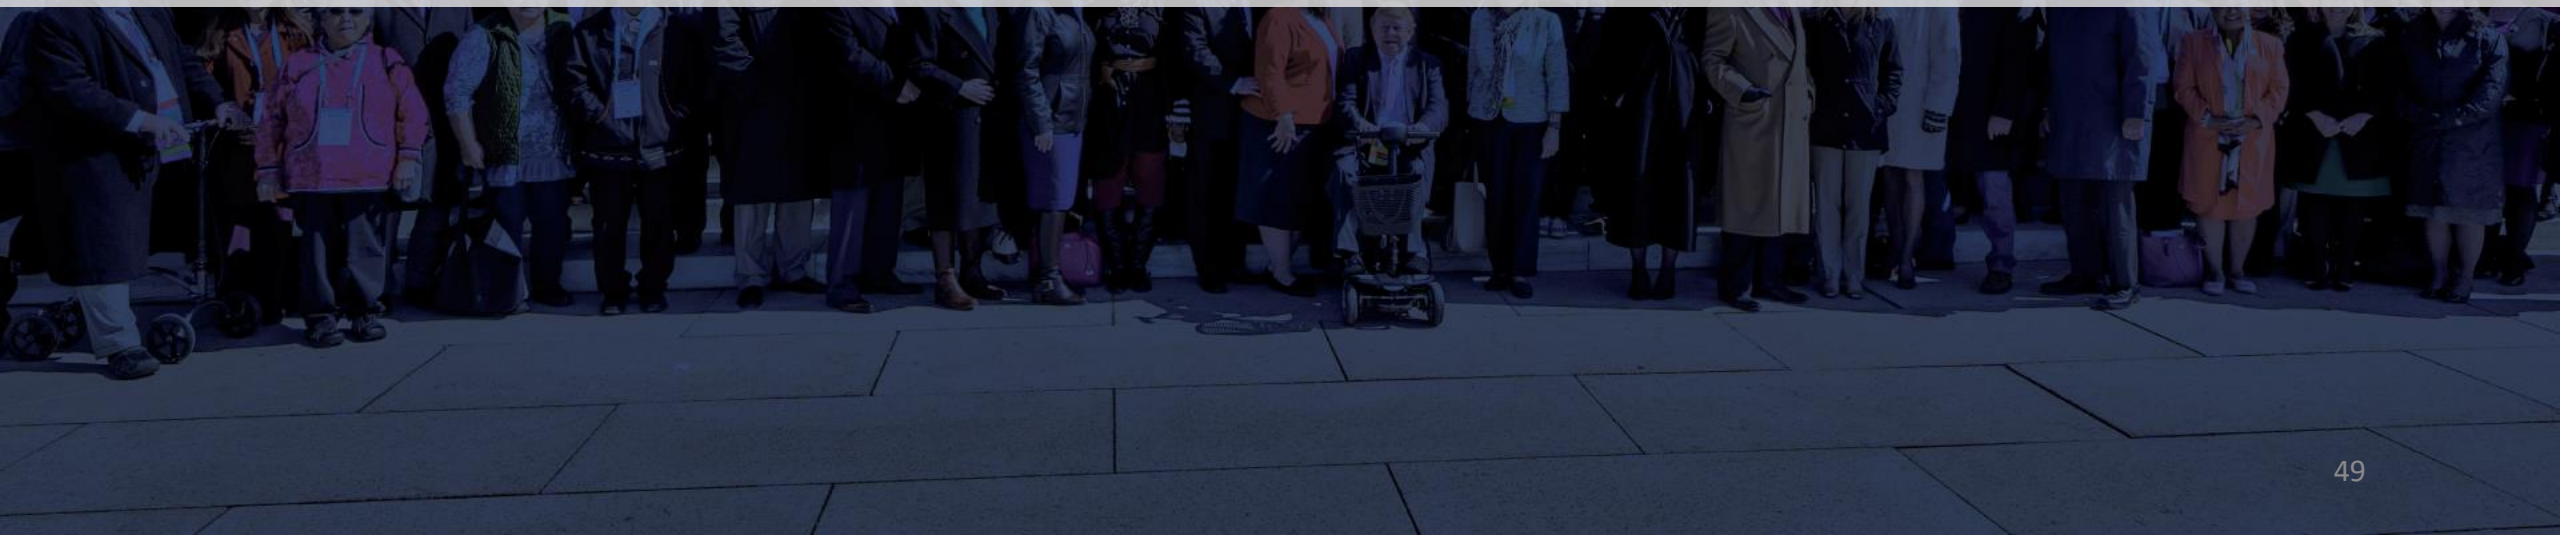

Status Report: 2025 NACo AI Policy Priorities

KEY ARTIFICIAL INTELLIGENCE POLICY PRIORITIES FOR COUNTIES

The rise of generative artificial intelligence (AI) has presented novel opportunities and challenges for the public and private sector alike. The current regulatory and legislative framework surrounding AI and generative AI presents opportunities for passing meaningful laws that will promote intergovernmental collaboration in a manner that will seek to protect human rights, monitor for the safe and responsible application of AI, and safeguard against nefarious uses of technology. State and local governments have already begun implementing AI to automate services in recent years, and as technological developments in generative AI continue to evolve, it will become necessary for new policy principles and practices to emerge in order to minimize the harmful impact that this technology could pose to society.

This analysis provides an overview of NACo's key 2025 legislative priorities for AI, including standards and guardrails to ensure that AI continues to bring meaningful innovation to counties and the greater society.

KEY POLICY HIGHLIGHTS

- Establish an **intergovernmental governance structure** that addresses the various uses of generative AI across different sectors.
- Dedicate a new **information sharing analysis center** (ISAC) for the creation of resource hubs and task forces, and development of an ongoing communication channel for intergovernmental coordination.
- **Provide direct funding assistance** to promote digital literacy and best practices, assistance for counties and workforce development.
- **Dedicate support mechanisms** to federal and local government agencies promoting the use of AI for public services.
- Mitigate the negative uses of generative AI in the **elections space.**
- Strengthen funding resources and regulatory oversight at independent agencies such to **combat mis- and dis- information geared towards consumers.**
- Implement federal guidance clarifying that **liability for outputs causing discrimination** rests with the owners and operators of AI models.
- Adopt and disseminate **data privacy** governance standards and best practices across all levels of government.
- **Require public engagement and participation** in AI policy-making processes to ensure the voices of diverse stakeholders are heard and considered.

KEY POLICY HIGHLIGHTS

- Establish an **intergovernmental governance structure** that addresses the various uses of generative AI across different sectors.
- Dedicate a new **information sharing analysis center** (ISAC) for the creation of resource hubs and task forces, and development of an ongoing communication channel for intergovernmental coordination.
- **Provide direct funding assistance** to promote digital literacy and best practices, assistance for counties and workforce development.
- **Dedicate support mechanisms** to federal and local government agencies promoting the use of AI for public services.
- Mitigate the negative uses of generative AI in the **elections space.**
- Strengthen funding resources and regulatory oversight at independent agencies such to **combat mis- and dis- information geared towards consumers.**
- Implement federal guidance clarifying that **liability for outputs causing discrimination** rests with the owners and operators of AI models.
- Adopt and disseminate **data privacy** governance standards and best practices across all levels of government.
- **Require public engagement and participation** in AI policy-making processes to ensure the voices of diverse stakeholders are heard and considered.

Following the Administration: White House AI Action Plan

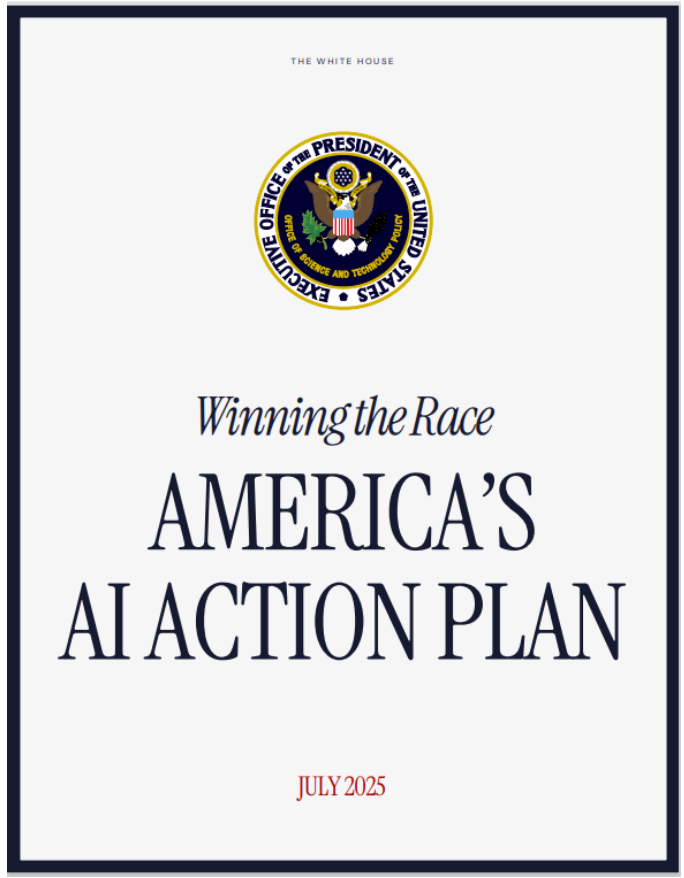

On July 23rd, the White House released their AI Action Plan, outlining 90 policy proposals on AI to federal agencies.

- Directives invite agencies to begin public rulemakings and internal initiatives to carry out the goals of the action plan.
- Action Plan includes a new AI-ISAC and directives to support workforce-targeted AI education and training.
- Counties should continue to monitor updates on AI from federal agency partners

Read NACo's Blog on the AI Action Plan Here

Image source: DLA Piper, May 2025.

Following Congressional Action on AI

Problem: In May, the U.S. House of Representatives introduced a provisions that would enact a10-year moratorium on state and local AI policymaking.

Advocacy: As the measure gained traction and passed the House, NACo conducted advocacy alongside key stakeholders to defeat this proposal in the Senate.

Result: The Senate ultimately pulled the provision from the reconciliation bill text by a vote of 99-1.

Image source: DLA Piper, May 2025.


Questions?

Artificial Intelligence County Compass: Practical Toolkit for Local Implementation

Metropolitan Washington Council of Governments
September 10, 2025

Rita Reynolds, CIO
National Association of Counties
rreynolds@naco.org

NACo's Journey

Spring 2023 – Awareness

May 2023 – July 2024 AI Exploratory Committee

July 2024 – July 2025 AI Regional Forums and Presentations

July 2025 – AI in Motion Use Case Resource

Aug – Dec 2025 – AI Regional Forums and Education

Policy In Action

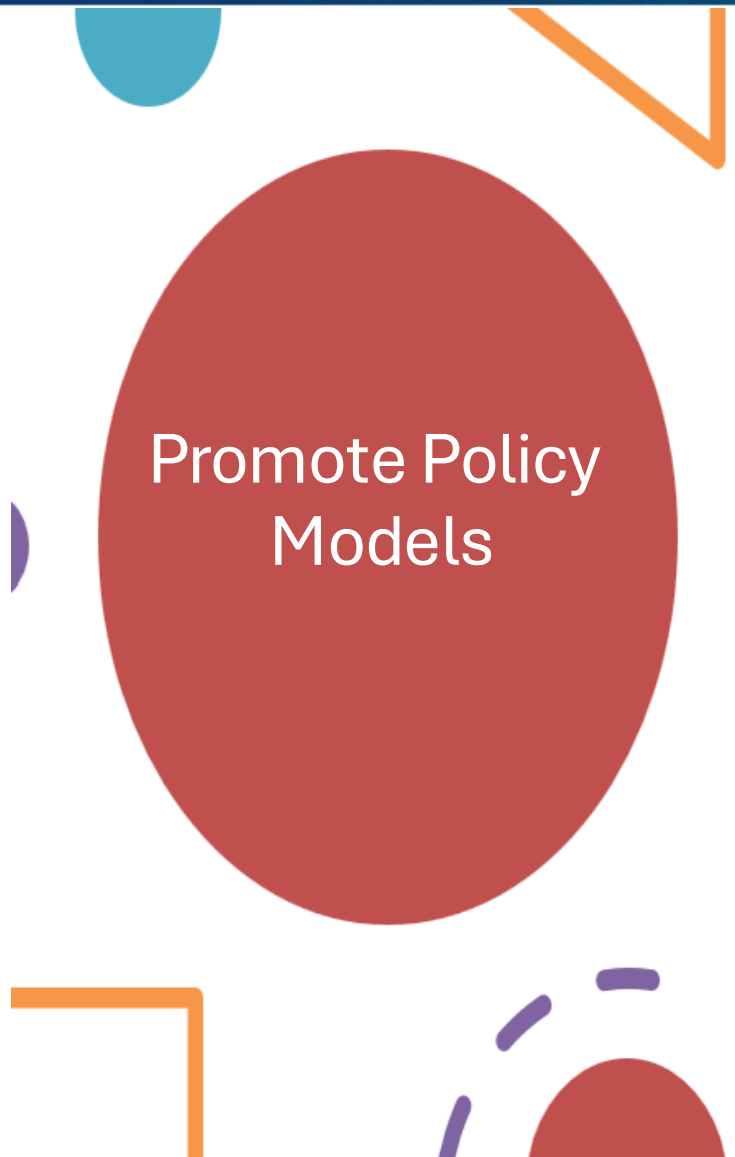

- Policy Framework:
  - Establish policy framework for GenAI
  - Review key legal considerations
  - Review and assess existing procurement policies

Policy In Action

Santa Cruz County Calif.

Approved 09/1
Board of Super
DOC-2023-769

County of Santa Cruz Artificial Intelligence Appropriate Use Policy

Purpose: The purpose of this policy is to establish County of Santa Cruz employee practices for the responsible and secure use of generative artificial intelligence (AI). The County is committed to utilizing Artificial Intelligence (AI) technologies responsibly and ethically to improve processes, enhance services to County residents, and support employees to do their best work. This AI Appropriate Use Policy provides simple, user-centric guidance for all employees, regardless of technical expertise.

AI tools are developing at an exponential rate. The County will regularly review and update this policy to keep it aligned with ethical and legal standards and technological advancements in generative AI as frequently as needed.

Scope: This policy applies to all employees, contractors, and any other third-party individuals or entities who have access to generative AI technologies or are involved in using generative AI tools or platforms on behalf of our organization.

1. Defining Artificial Intelligence (AI): For the purposes of this policy, Artificial Intelligence, also known as machine intelligence, is the simulation of human intelligence processes, such as problem solving by machines.

2. Defining Generative AI Tools: Generative AI tools are computer programs capable of many activities, including but not limited to completing general administrative office tasks, data analysis, programming, and image creation. While these tools can improve productivity, it is crucial to use them responsibly to comply with various laws, maintain data privacy and security, and uphold County values.

I. Principles for Responsible Generative AI Tools Use: Staff should be open to responsibly incorporating Generative AI into their work where it can be beneficial for making services better, more just, and more efficient. Each employee is responsible for using generative AI tools in a manner that ensures the security of sensitive information and aligns with County policies. Here are key principles to follow:

Ethics In Action

Establish an Ethical Framework (Transparency)

Keep the Human in the Loop

Foundational ethical principles for use of GenAI should include:

Fairness, Equitableness, and Impartiality

Transparency

Privacy

Accountability

Applications In Action

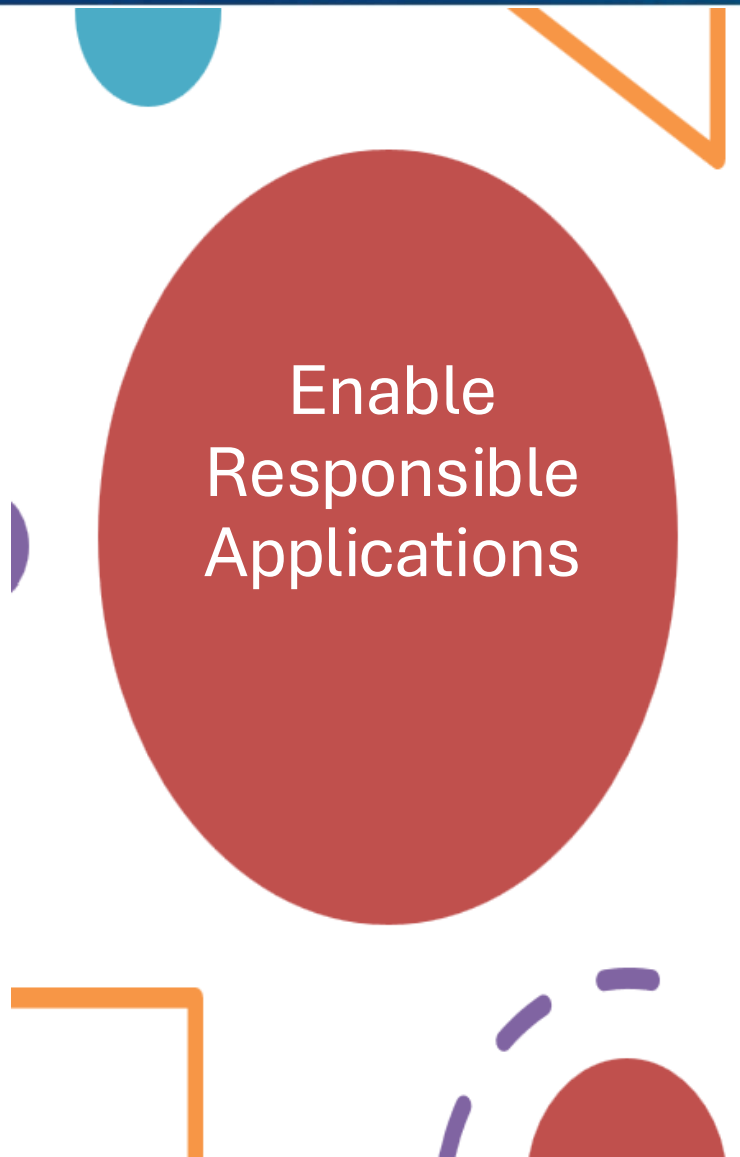

- Applications Framework:
  - **Review and evaluate use cases**
  - Familiarize with federal resources
  - **Practice robust data governance**
  - Regularly assess resources
  - Update cybersecurity measures
  - Design procedures for data training
  - Determine software, hardware, and procurement standards

Applications In Action

Evaluate Use Cases & Practice Robust Data Governance [European guidance]

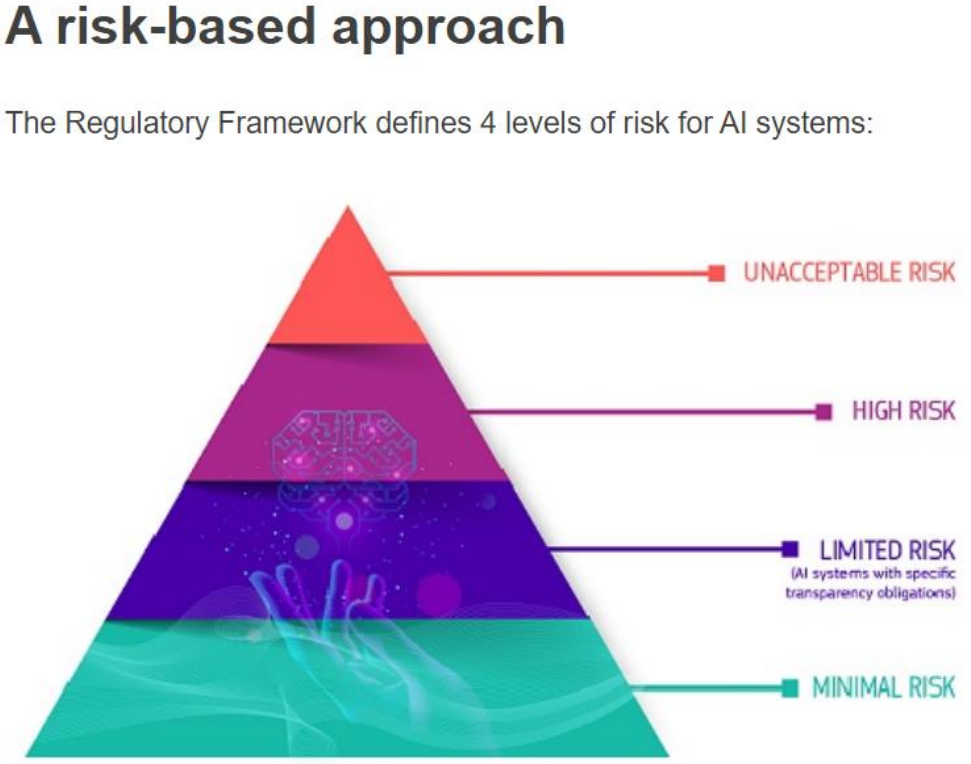

Examples: Low Risk – Press Release
High Risk – Mental Health Evaluations

Workforce In Action:

Preparing the Workforce

- Workforce Preparation:
  - Focus on skills development and training
  - Consider skills acquisition options
  - Develop a multi-year workforce strategy
  - Inform and seek feedback from workforce

AI is not going to replace humans, but humans with AI are going to replace humans without AI

Harvard Business Review

Challenges & Benefits

Governance and compliance

Security and privacy

Copyright issues

Accuracy validation

Preventing bias and ethical issues

Managing change and trust

Training county staff

General Productivity

Optimize social services

Improve public safety

Personalize service delivery

Create tailored local solutions

Utilize forecasting

Engage community stakeholders

NACo AI in Motion Web Resource

The representative counties come from a cross section of states and can be categorized into five themes

Government Operations & Workflow Automation
Streamlining internal processes, automating manual tasks, and improving productivity.

Public Service & Resident Engagement
Enhancing public access to information, legal assistance, and service delivery.

Emergency & Resource Planning
Using historical data and AI to forecast demand and enhance emergency response readiness.

Education & Population Planning
Forecasting demographic changes and infrastructure needs using AI.

Cybersecurity & Data Privacy
Enabling secure, controlled use of AI within government environments.

Take aways

Clean up your data!

- ***Educate*** *(webinars, in-person events, virtual trainings)*
- ***Assess*** *staff utilization (inventory)*
- ***Ideate*** *(Tabletops, AI hackathons)*
- ***Conduct Pilot*** *(low-risk, productivity areas)*

AND

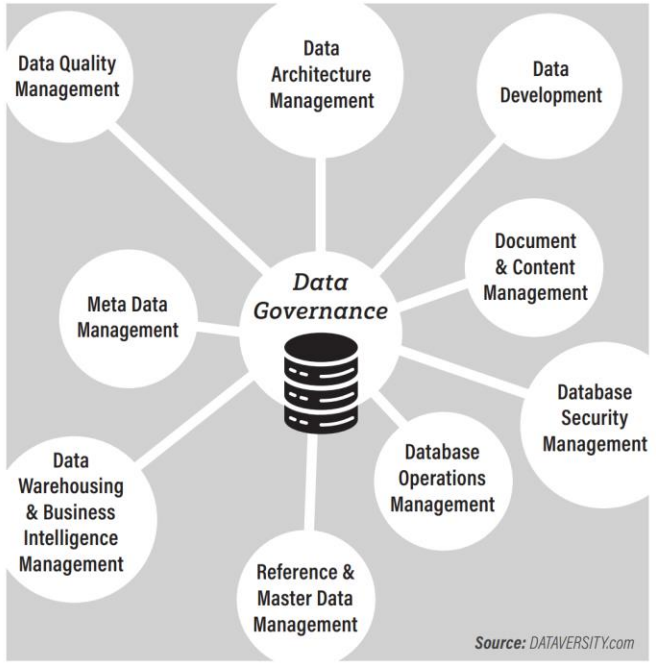

Metropolitan Washington
Council of Governments

AI POLICY IN ACTION

In 2024, lawmakers introduced
150+ AI bills nationwide

30+ states issued AI policy or guidance

(NCSL 2025)

Houston-Galveston Area Council

AI Activation Event

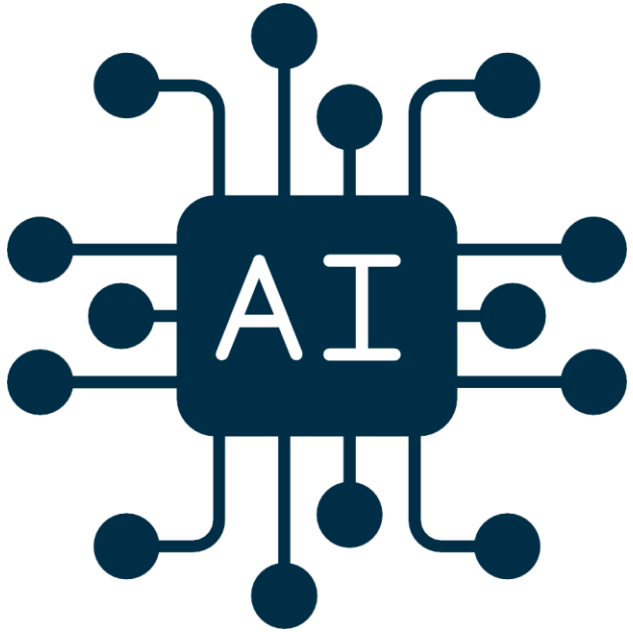

Welcome to the AI Parking Lot

Houston-Galveston Area Council

Roadmap

Phase I: Intro to AI
May 2025

Phase II: AI for Our Work
June 2025

Phase III: Internal Chatbots
End of 2025

Phase IV: Document Optimization
Late Spring 2026

Phase V: External Chatbots
Summer 2026

Phase VI: Data Prediction
End of 2026

Questions? Contact the Data Services team anytime.

AI Parking Lot

Houston-Galveston Area Council

Your AI Parking Lot at a Glance

TOOL	BEST FOR	KEY FEATURES
ChatGPT	Writing, summarizing, asking questions, brainstorming	- Works with text *and* images - Easy to adjust tone - Multilingual
Llama	Quick tasks, checklists, translations, simple writing	- Fast, text-only tool - Lightweight and efficient - Supports 8 languages
Claude	Editing, outlining, creative ideas, longer content	- Strong at rewriting & feedback - Great for structure & clarity - Most natural tone

Each AI Model can: Search the web, upload and analyze files, & summarize websites and research

METROPOLITAN WASHINGTON COUNCIL OF GOVERNMENT

BOARD OF DIRECTORS

Rodney Lusk, *COG Board Chair*
Clark Mercer, *COG Executive Director*

Metropolitan Washington
Council of Governments

REGIONAL BEST PRACTICES

LEARNING FROM LOCAL LEADERS

Bryan Hill
Fairfax County Executive
(moderator)

Stephen Miller
CTO
District of Columbia

Vanetta Pledger
CIO
City of Alexandria

Gail Roper
CIO
Montgomery County

Nate Wentland
CIO
Loudoun County

Metropolitan Washington
Council of Governments

BRINGING IT ALL *TOGETHER*

Metropolitan Washington
Council of Governments

PULSE SURVEY RESULTS

WHERE OUR ORGS ARE ON AI TODAY

- **ADOPTION IS BROADENING.**
  - 25% report multi-team AI use; another 48% are in pilots or individual-only use.
- **GOVERNANCE IS FORMING.**
  - 29% have an approved AI policy; 42% are draft or unsure.
- **PATH TO SCALE.**
  - The biggest lift is converting pilots and individual use into sanctioned team workflows; formalize guardrails to unlock momentum.

Metropolitan Washington
Council of Governments

AMAZON WEB SERVICES

About the Founder

Father. Husband. Leader. Author

- Founder & CEO of **OptimaNova AI**
- Chief Executive Officer of **Byte Back**
- 29th Executive Director of **Alpha Phi Alpha**
- COO of **The Stafford Foundation**
- Founder & CEO, **Campus 2 Careers**
- Human Resources Officer, **DC Government**
- DRA, **Management Leadership for Tomorrow**
- HR Manager, **Save the Children**
- FSU National Board of Directors
- Black Men Vote Board of Directors
- Leadership Greater Washington
- Author
  - "Morning Cup of Joe"
  - "100 Ways to Change the World"
  - "AI for Good"

TECHNOLOGY

STARTUPS TO WATCH

These 11 upstarts are transforming their industries and making a mark on D.C.'s innovation scene.

WASHINGTON BUSINESS JOURNAL

OptimaNova AI

About the company: Amid a market flooded with hundreds, if not thousands, of commercially available AI tools, OptimaNova AI LLC hopes to match companies with the ones that are most suitable to maximize the technology. Joe Paul, the former CEO of free computer training and certification provider Byte Back, founded the Northeast D.C.-based company in November 2023 to prepare and train organizations on how AI technology can best be used across a business. OptimaNova determines this by having its customers complete a free assessment test to shed insights on the types of commercial or bespoke AI tools that can best aid a company. The startup then offers consultation and training services on these tools to find ways to implement the technology directly for its customers or partners, which include consumer goods giant Procter & Gamble Co. and D.C. marketing firm The Brand Guild. OptimaNova has 23 employees in full-time and advisory-related roles and has not yet raised any outside funding.

Why we're watching: Over the next year, OptimaNova plans to launch new AI-based products specifically tailored to nonprofits and government organizations. It's also eyeing its first possible outside investment opportunity to help reach its next phase of growth. OptimaNova is on track to finish the year with over $1 million in revenue, which could reach $1.5 million depending on the outcome of some of its government contracting awards. – *Nate Doughty*

Joe Paul is the founder of D.C.-based OptimaNova AI.
JOE PAUL

OptimaNova AI

OPTIMANOVA AI
www.optimanova.ai

Leveraging AI to build

FUTURE READY ORGANIZATIONS

FULL-SCALE AI IMPLEMENTATION

OPTIMANOVA AI

What We Do

We make AI adoption effortless through end-to-end implementation

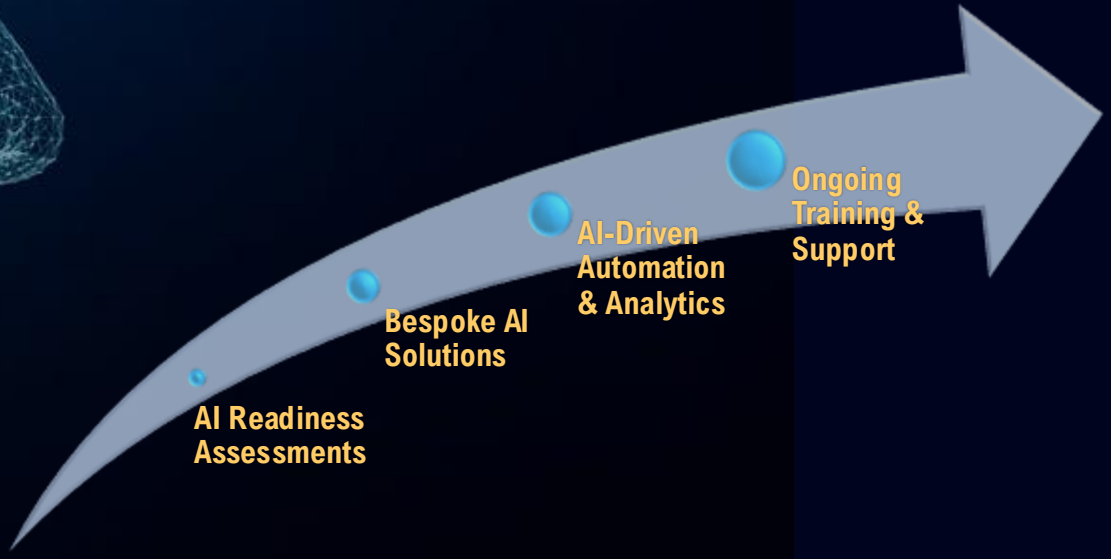

OPTIMANOVA'S AI SOLUTIONS

myvelocity.ai

Run your company like a Fortune 500 company with AI at the wheel.
Velocity centralizes nonprofit operations, automates grant writing, and visualizes impact, freeing teams to secure more funding.

Ideal for c3 organizations and schools seeking sustainable funding and operational velocity.

LEARN MORE ABOUT VELOCITY

OPTIMANOVA AI

www.optimanova.ai

MEET OUR DYNAMIC AI AGENTS

"NOVA"

AI Business Strategist

Automates operations, boosts efficiency, and drives AI-powered business growth.

"AMARA"

AI Policy & Ethics Advisor

Ensures AI governance, ethical compliance, and responsible implementation.

"APOLLO"

AI Technical Architect

Designs scalable AI systems for seamless tech integration and innovation.

"ATLAS"

AI Compliance & Risk Analyst

Protects data, enhances AI security, and ensures regulatory compliance.

"IVY"

AI Learning & Education Expert

Advances AI-driven training, upskilling teams for the future of work.

"ORION"

AI Research & Innovation Assistant

Explores AI trends, forecasts industry shifts, and drives innovation.

ADONIS' ADVENTURES

- **AI now touches** benefits, policing, health, housing, jobs; every model is a policy decision in code.
- **Equity is not charity**; it is accuracy, legality, and public trust.
- **North Star**
  - Build AI that sees every resident clearly, including the kid who thinks he can fly.

FOUR FACTS YOU CANNOT UNSEE

HEALTH CARE

Algorithm gave less care to Black patients. Fix = referrals jump **17.7%** → **46.5%**.

SPEECH-TO-TEXT

Error rate nearly **2x higher** for Black speakers. **23% unusable** vs. 1.6% for white.

FACE RECOGNITION

False positives **10–100x higher** for Africans, Asians, women. Errors flip ID outcomes.

MEDICAL IMAGING

AI reads race from X-rays (**AUC 0.91–0.99**)—signal humans can't see, bias baked in.

WHAT THIS MEANS FOR MWCOG

RISK

If we do nothing, inequity scales at machine speed and erodes trust in digital services.

GUARDRAILS

We already have the **AI Bill of Rights** and **NIST AI Risk Management Framework**—local governments can use them today. *(White House, NIST)*

PLAYBOOK

Counties have **NACo's AI County Compass** for risk tiers, workforce prep, and policy models. Use it. *(NACo)*

CALL TO ACTION

BUILD EQUITABLE AI IN OUR REGION

- **ADOPT GUARDRAILS NOW**
  - Use the NIST AI Risk Management Framework and AI Bill of Rights as baselines.
- **TEST WHAT WE BUY**
  - Require bias and equity audits for every AI procurement; no test, no deploy.
- **CREATE LOCAL BENCHMARKS**
  - Stand up shared tests for 311, 911, benefits, and translation systems using real regional voices.
- **ENGAGE OUR RESIDENTS**
  - Bring families, schools, and communities into AI literacy so trust grows alongside technology.