Google for Government

Building a Secure Foundation for American Leadership in AI

Karen Dahut,
CEO, Google Public Sector

Katharyn White,
Director of Marketing, Google Public Sector

Phil Venables,
CISO, Google Cloud

Jacob Crisp,
Global Head of Strategic Response, Google Cloud

Introduction

NY Times columnist and author Thomas Friedman recently argued that humanity faces a Promethean moment with the advent of artificial intelligence (AI). Just as the Prometheus of Greek mythology is said to have bestowed the gift of fire upon humanity and gave rise to civilization, AI has the potential to usher in a new era in human progress, similar to the invention of the printing press, and the scientific and industrial revolutions.

At Google Public Sector, we believe that advances in AI will drive the biggest technological shift we will see in our lifetimes, and we're excited about the opportunities that lie ahead. As Google's CEO recently wrote, it's bigger than the shift from desktop computing to mobile, and it may be bigger than the internet itself. It's a fundamental rewiring of technology and an incredible accelerant of human ingenuity. AI is already having a positive impact on the way we live and work, for example, and how governments deliver services, including enabling online translation across dozens of languages and speeding up support for humanitarian relief after natural disasters.

To help advance this future, we're committed to making AI more helpful for everyone, and deploying it responsibly. But governments will also play a central role in helping manage this transformation. The United States, for example, drove whole-of-society approaches to global projects like the Apollo space program and the global Internet (which led to modern benefits for everyone). When it comes to AI, we need a similar–perhaps even more expansive–effort today. That's why we've supported the White House's effort to work with industry to implement voluntary AI commitments, and we've since provided our first public update on our progress. This builds on our enduring priority to invest in and grow the nation's cybersecurity workforce and AI expertise and skills specifically. Here are a few other ways we've helped:

Launched a new cybersecurity certificate and gen AI skills and training program

Committed to train 100,000 Americans in fields like data privacy and security

Announced a new research program for learning and career opportunities

Expanded access to cybersecurity clinics

Introduction

As we build and deploy this technology together, one of the most pressing challenges is to ensure that it is protected from digital threats. There are three important drivers that create an urgent need to address this challenge now. First, the impact of AI on modern society is so profound and advances are happening so fast that the United States must take immediate steps to stay ahead of potential threats. Second, we know that it's important for technology providers to think about security ahead of time, not after an attack has occurred. Indeed, we must work together now to build a more secure foundation for America's future. Third, it is clear that advances in AI will transform the future of work, which creates a growing mandate for change and perhaps the biggest opportunity in generations for improved leadership and growth.

It is absolutely critical that society gets this right. As AI continues to develop, the intersection of AI and security will become a critical part of every organization's AI journey. Not only will getting AI security right matter for America's ability to protect our data, our critical infrastructure, our nation, and ourselves, but also it will define American AI leadership for decades to come. To be successful, we need a holistic, ecosystem-wide approach (e.g., data, models, hardware, applications, supply chain, among many other components) to address fundamental issues and challenges systemically. This is consistent with Google Cloud's shared fate model with customers, and it creates a shared responsibility when it comes to AI - to train and tune the models, to verify and check outputs for accuracy, and to be transparent about how we use these technologies. That's what it means to be bold and responsible in this space.

Google pledges $10 billion for cybersecurity over five years at White House Summit in 2021

Introduction

We've been working at this intersection for over a decade (including years of foundational AI research by Google and DeepMind and AI-powered security innovations in Workspace and Safe Browsing, to name a few) and we've evolved our approach to include building products while using secure-by-default and secure-by-design principles. We approach AI systems in the same way we view other security challenges: we bake in industry-leading security features (often invisible to users) and robust protections to keep our users safe. We look at this investment like a digital immune system — when we learn and adapt from previous risks to our digital health, our systems become better equipped to protect against, anticipate, and predict future attacks.

Google's approach is critical for the public sector, where we have the largest cloud service offering available on the market for U.S. public sector customers. Our software-defined community cloud (click here for more information on a new way to "Government Cloud") allows customers to take advantage of Google Cloud's efficient cloud infrastructure to help support stringent security compliance requirements and the latest innovation without needing to manage and maintain individual infrastructure instances. Google was a pioneer in developing and implementing a Zero Trust approach to security—the idea that every network, device, person, and service is untrusted until it proves itself—and has been helping global governments and organizations adopt Zero Trust architectures. Google Cloud also relies on defense in depth, with multiple layers of controls and capabilities to protect against the impact of configuration errors and attacks.

These capabilities, if implemented properly and at scale, can help organizations maximize AI effectiveness. This includes helping to address emerging regulatory requirements and AI model and data governance. More importantly, public sector leaders can leverage AI to create immediate value for constituents and society as a whole, particularly in cases that do not require policy changes or new sources of data.

In recent engagements with customers, we are focused on three categories of potential use cases:

01

Improving constituent, customer, and end user outcomes

02

Increasing government employees' efficiency and effectiveness

03

Transforming process or operations

What's possible with a secure AI foundation

Google has been working with the U.S. public sector to execute on the needs of the mission, using data to improve the constituent experience and offer more accessible digital services, while also helping government workers operate more efficiently. AI is a critical part of that effort, and building and deploying a secure foundation to deliver these technologies is essential to helping agencies achieve these goals.

AI can improve the efficacy of interacting with government and increasing trust, in part by offering even more insightful and natural conversations in real time and reducing time for constituents to find the information they need. Government entities can use AI to extract information and help automate paper-based processing, and create systems that automatically generate summaries and reports on government activities. These reports could help increase transparency and accountability, and improve trust around how to best allocate resources. Further, AI can help enhance efficiencies in operations by identifying patterns in data, predicting trends and identifying fraud, removing variability, and making recommendations.

We're already starting to see progress on these fronts, and we're excited about the potential for AI to unlock the **following possibilities:**

Government becomes the most efficient service provider in the world

Departments and agencies are integrating critical datasets to make more informed decisions on how to best support their communities, and AI can help communities get better access to government services. Over time, we believe governments can become the most efficient service providers in the world.

The City of Memphis, for example, worked with Google Public Sector and Egen (formerly SpringML), a Google Cloud Premier Partner, to find and fix potholes and identify vacant properties through the use of AI and Machine Learning (ML). The team projects 75% of potholes will be identified as a result, improving streets, communities, and the experiences of both residents and visitors.

Google Public Sector is also working with the city of Dearborn, Michigan, where more than half the population speaks a language other than English at home, to deploy AI to the city website. That way, constituents who speak Arabic and Spanish can understand and interact with their city and make use of its services.

Government helps transform personalized education and deliver next generation workforce

AI offers benefits in the education space, which faces constant demands for scale, personalization, and innovation to prepare the next generation for success. It can improve user experience for students, staff, and faculty, acting as an assistant, and provide the workforce with opportunities to strengthen their skills and stay up to date with training. It can also accelerate content development and customize large models to analyze the ever-growing mountains of data generated across organizations everywhere.

Educators, researchers, IT professionals, student developers, and C-suite leaders are starting to explore AI and gain insights into how it is already transforming teaching, learning, and advanced research across the country. Governments play an important role in this effort.

Google customers Form Bio and Collegis Education, for example, are using AI to make research and student services more efficient. By analyzing how and when students, faculty, and staff interact with online services, Collegis can determine the best times and methods to reach each cohort. Form Bio builds AI into their scientific data platform to help researchers focus on their results. Making computational workflows user-friendly can save time and resources–and accelerate discoveries.

Government helps enable better healthcare and reshapes training and skills programs

Healthcare breakthroughs change the world and bring hope to humanity through scientific rigor, human insight, and compassion. We believe AI can contribute to this, with thoughtful collaboration between researchers, healthcare organizations, and the broader ecosystem.

In September, the Defense Innovation Unit (DIU) published information on the Department of Defense's efforts to prototype and field AI solutions that will help medical experts transform and improve health care. Google Public Sector worked with DIU and others to build an AI-powered microscope that can help doctors identify cancer. The collaboration gives pathologists an AI assistant that helps them deliver more accurate and timely cancer diagnosis, transforming the healthcare experience for the military community and beyond.

Government revolutionizes emergency services and public safety response

Between January 1 and September 11, 2023, there have been 23 confirmed weather/climate disaster events in the United States, causing more than $1 billion in losses. These disasters resulted in 253 deaths. Faced with rising challenges and limited resources, departments and agencies need better capabilities to help accelerate federal emergency assistance and disaster recovery efforts. AI can offer a range of benefits, including drawing insights from disparate data sources and helping enhance efficiency, accuracy and responsiveness.

Charged with serving as the federal lifeline for millions of citizens who need immediate help in the face of life-threatening disasters, the Federal Emergency Management Agency (FEMA) is working with Google Public Sector to run its data management system more efficiently, securely, and collaboratively. Earlier this year, Sam Hultzman, FEMA's National Flood Insurance Program's Pivot System Owner, shared how the agency moved from a mainframe to Google Cloud, and in the process built a "one-stop shop" for all their customers – which includes states, counties, insurance companies, internal users, and data scientists – to access the applications they need.

Government fuels cutting edge R&D and scientific innovation to advance human evolution

AI has the power to transform the way R&D teams work and the manner in which research and development teams operate, as well as fuel innovation. It allows scientists to accelerate their discovery process, streamline their research, and gain new insights into their products and consumers.

The National Ecological Observatory Network (NEON), the National Institutes of Health (NIH) STRIDES program and NCI Imaging Data Commons, for example, are using Google Cloud to help accelerate research productivity with purpose built, scalable data management tools. These organizations collect large amounts of data from disparate sources, which need to be easily accessible to researchers around the globe to advance important scientific discoveries. With non-restrictive, open-source technology, Google Cloud offers a robust infrastructure with data management tools for the petabytes of data collected by the institutions.

How to get started building a secure AI future today

One of the benefits of our experience using AI to solve real-world problems is that we have become better at helping secure new technologies as they become mainstream. At the same time, we're leveraging recent AI advances to provide unique, up-to-date, and actionable threat intelligence, improving visibility across attack surfaces and infrastructure.

We know that improving cybersecurity is no longer a human-scale problem, and we're excited about continuing to work with governments to meet tomorrow's security threats.

To help organizations get started, we're sharing three key building blocks for a holistic, ecosystem-wide approach to AI and security. If implemented properly, these blocks will form the basis for a secure foundation for American leadership in AI, and help maximize the benefits of AI technologies and minimize risks. One of the common themes across these threads is the need for robust collaboration between the public and private sectors on shared opportunities, threats and challenges. Google remains deeply committed to supporting those collaborative efforts.

At Google Cloud, we've built our approach to data governance and responsible AI into Vertex AI, our AI platform for building and deploying ML models. Vertex AI enables enterprises to use their own data to tune generative models for their use case, while maintaining control over their data and IP. To help protect this data, we offer a number of security and compliance controls for generative AI, including VPC Service Controls for mitigating the risk of data exfiltration from Vertex AI, and customer managed encryption keys, access transparency, and data residency. Google Cloud is committed to the same data responsibilities for ML data that we apply to other customer data.

Building Block 1

Understand threat actor interest in, and use of, AI capabilities

In the first half of 2023, we saw a 35% increase in incident response engagements compared to the same period last year. Through this ongoing, frontline engagement, our experts saw four key trends:

1. All apex threat actor groups are now active at the same time;
2. Cybersecurity is now a top geopolitical priority for most governments;
3. Threat actor groups are now professionalizing operations and programs;
4. Threat actor groups' tactics now evade "standard" controls.

Google Cybersecurity Clinic

We've also observed unprecedented developments like the first time cyber operations played a prominent role in war. The threat landscape remains dynamic and complex, and we expect these trends to continue in 2024 and beyond.

These trends are equally important in the context of AI. Since at least 2019, we've tracked threat actor interest in, and use of, AI capabilities to facilitate a variety of malicious activity. Based on our own observations and open source accounts, adoption of AI in intrusion operations remains limited and primarily related to social engineering. In contrast, information operations actors of diverse motivations and capabilities have increasingly leveraged AI-generated content–particularly imagery and video–in their campaigns, at least in part because of the readily apparent application of AI to disinformation.

Every department and agency must become deeply familiar with the current threat landscape, and public sector leaders at all levels of government need to understand the connection between threat intelligence and risk mitigation. In most cases, cybersecurity leaders understand the need for better intelligence on threat actors, but many of them make decisions without fully understanding who may seek to target their organization, why, and how. In a recent survey of business and IT leaders, more than three-quarters of respondents said that they make decisions without insights on who could be targeting

their organization and only about one-third had a comprehensive understanding of different threat groups and their tactics, techniques, and procedures. These visibility gaps mean defenses may not meet their intended goals. Organizations can work to bridge these intelligence gaps and ensure this information is playing a leading role in risk management decisions.

As AI technology evolves, we believe it has the potential to significantly augment malicious operations in the future, enabling threat actors with limited resources and capabilities, similar to the advantages provided by exploit frameworks including Metasploit or Cobalt Strike. As a result, we expect to see more adversary use of AI tools over time. However, we also believe governments can scale to meet these threats with strong threat intelligence programs and robust collaboration, in addition to using these same technologies to tip the scale towards defenders. We'll explore these capabilities and how they help government employees and constituents in more detail below.

Building Block 2

Deploy secure AI systems

As public sector leaders look to deploy AI technologies, they should seek to understand the risks associated with AI systems to protect their organizations.

It is important to note that securing AI does not mean you need to upend security best practices, and much of the wisdom that security teams have learned is still correct and applicable. But it does have some important differences and AI systems can introduce new security risks that are not present in traditional systems. For additional information on these risks, we recommend reviewing Google's AI Red Team paper.

To help governments navigate these complexities, we recently launched the Secure AI Framework (SAIF). SAIF is a conceptual framework for secure AI systems. It is inspired by security best practices—like reviewing, testing and controlling the supply chain—that Google has applied to software development, while incorporating our understanding of security mega-trends and risks specific to AI systems. We believe that every department and agency will need to pursue this or a similar approach to mitigate risks across their technology stack.

SAIF offers a practical approach to address the concerns that are top of mind for security and risk professionals, such as security, AI/ML model risk management, privacy and compliance, and people and organization.

To get started putting SAIF into practice, organizations should take four steps:

1. Understand the use
2. Assemble the team
3. Level set with an AI primer
4. Apply the six core elements of SAIF

For more information on these steps and how departments and agencies can implement SAIF

Click here

Six core elements of SAIF

1

Expand strong security foundations to the AI ecosystem

2

Extend detection and response to bring AI into an organization's threat universe

3

Automate defenses to keep pace with existing and new threats

4

Harmonize platform level controls to ensure consistent security across the organization

5

Adapt controls to adjust mitigations and create faster feedback loops for AI deployment

6

Contextualize AI system risks in surrounding business processes

---

We're already taking steps to support and advance the framework and look forward to working with departments and agencies on this effort. **This includes:**

01 Fostering industry support for SAIF through continued industry engagement to help develop the NIST AI Risk Management Framework and ISO/IEC 42001 AI Management System Standard, as well as contributing to planned updates to the NIST Cybersecurity Framework and ISO/IEC 27001 Security Management System;

02 Working directly with organizations, including customers and governments to help them understand how to assess AI security risks and mitigate them;

03 Sharing insights from Google's leading threat intelligence teams like Mandiant and TAG on cyber activity involving AI systems;

04 Expanding our bug hunters programs (including our Vulnerability Rewards Program) to reward and incentivize research around AI safety and security;

05 Continuing to deliver secure AI offerings with partners like GitLab and Cohesity, and further develop new capabilities to help customers build secure systems.

Building Block 3

Supercharge security with AI

As the threat landscape evolves, we believe that departments and agencies need earlier detection, more automation, and assistance and expertise from top defenders to counter adversary activity.

Recent advances in AI and Large Language Models (LLMs) accelerate our ability to help defenders who are responsible for keeping government employees, constituents, and users safe. These technologies are advancing rapidly, and governments must be familiar with what they offer as soon as possible and explore options to deploy them for mission success.

At Google, we are infusing AI across our security products. The Google Cloud Security AI Workbench is an industry-first extensible platform powered by a specialized, security LLM, Sec-PaLM 2. This model is fine-tuned for security use cases, incorporating our proprietary security intelligence on vulnerabilities, malware, IOCs, and threat actors.

These capabilities not only give government cybersecurity leaders a more natural and creative way to understand and manage their environments, they give them access to AI-powered expertise to go beyond what they could do alone. For example, whereas traditional scanning tools might tell you whether a file or URL is malicious, our VirusTotal Code Insights capability, powered by AI, can offer defenders natural language insights into the malware detected, how it works, and how to defend against it. Every department and agency should be thinking about how they can use AI to achieve better security outcomes.

To help proactively address security challenges for customers, we've recently expanded our AI capabilities with Gemini, our AI collaborator that provides generative AI-powered assistance to cloud defenders where and when they need it. And we continue to launch new security innovations and enhancements across our security operations and cloud platforms. In the same way that a consumer can now prompt a foundation model to summarize news items, these capabilities can help government security teams query their own organization's data for insights into cyber threats.

Conclusion

Like many of our customers, government organizations will continue to face security challenges, including threat overload, toilsome tools, and the talent gap. But based on over two decades building and deploying secure systems, we believe that advances in generative AI can help tackle these common challenges faced by security professionals - keeping pace with attacks, managing numerous tools, and upskilling staff to help address cyber risk.

Of course, AI is not a silver bullet and so, during times of great change, we believe in the importance of collaborative leadership. Our government must lead with strong technology partnerships and individuals who champion change. We're optimistic that our continued, shared work will build on the incredible progress we've already seen and continue to create a foundation for impactful leadership and the safe, secure, and trustworthy development of AI in the years to come.

Google for Government