Statement of Dr. Andrew Moore
Vice President and General Manager – Cloud AI & Industry Solutions, Google Cloud
Senate Armed Services - Cybersecurity Subcommittee
Tuesday, May 3, 2022

Chairman Manchin, Ranking Member Rounds, and members of the Committee, thank you for the opportunity to appear before you this morning.

My name is Andrew Moore. I am Vice President and General Manager of Google Cloud Artificial Intelligence (AI). I most recently served as a Commissioner on the National Security Commission on AI (NSCAI) and I currently serve as a task force member on the National AI Research Resource (NAIRR). I previously served as Dean of the Carnegie Mellon University School of Computer Science and have spent my career as a computer scientist specializing in machine learning and robotics. I have also spent time as an advisor to the Department of Defense as a member of Google Cloud's leadership team.

I appreciate the Committee's support for advancing AI — thank you Chairman Manchin for your leadership in driving a partnership between the National Science Foundation and West Virginia University to ensure more funding for AI research in last year's appropriations bill, and thank you Ranking Member Rounds for your continued support of AI baselining at the Department of Defense. And I greatly appreciate the support both of you have provided for the NSCAI and its work. During my time there, NSCAI submitted strong recommendations to the Committee and the Department of Defense (DoD). In addition to the NSCAI recommendations, it is also worth revisiting the recommendations led by the National Academies and sponsored by the Office of the Director of National Intelligence on the Implications of Artificial Intelligence for Cybersecurity. AI can be an incredible asset but, as with any new technology, can also present new vulnerabilities.

A useful definition of AI is a machine which seems to have human or sometimes superhuman capabilities at a task we might previously have said needs uniquely human intelligence. In recent years some of the biggest advances have come from neural networks, which simulate billions of neural connections in biological nervous systems. The two big technological battles happening in academia and corporations around the world are first, how to scale it up to trillions of connections, and second how to turn really amazing technology demonstrations into practical deployed systems that are actually useful.

AI can refer to any number of technologies involving artificial systems designed to or having the ability to learn. One way Congress has itself described it is as "An artificial system designed to think or act like a human, including cognitive architectures and neural networks." A neural network is a computation system used to classify and analyze data using a process that mimics the function of the human brain. The data is fed into the first layer of a neural network, with each layer making a decision, then passing that information onto multiple nodes in the next layer. Some modern neural networks have hundreds or thousands of layers, with millions and even billions of parameters – the output of which can do such things as classify an object, or find patterns in data. This means that AI can process more information more quickly than a human: finding patterns and discovering relationships in data that any human would never be able to process on its own given the volume data being processed. And, AI is not limited by time of day, the need for breaks, or other human encumbrances. In the cloud, AI and machine learning can be "always on," continuously working on their assigned tasks.

For cybersecurity and in the context of national security, having the upper hand in AI against your adversary is critical. There is a race to see who can get machines to provide as much defense as possible. For example, AI systems are absolutely necessary to automate aspects of cybersecurity. The US remains the leader in AI, but we must ensure we continue to do this at scale.

AI powers all Google's products. And, importantly, we use AI to monitor our network infrastructure and attempt to predict and detect threats to our network or users. One of AI's critical uses is finding anomalies in activity that would indicate a new threat vector.

We of course also use it to support users when you search using Google Search. AI enables the most relevant responses to surface. AI is used to help predict the best route for you in Google Maps, detect misspellings or grammar mistakes in Google Docs and more. AI makes our products better by making them work for the user, by understanding and anticipating the user's preferences and needs. The same AI technology is used at Google to keep our users secure from phishing attacks on email, from malicious actors hacking into documents, and more.

AI also powers a lot of the solutions Google uses to serve the Department of Defense. For example, one of my favorite partnerships between the Department and Google Cloud is with the US Navy, where commercial drones are used by the Navy to take millions of images of the hull of ships and other hard-to-reach parts of ships, and then sends the images to Google Cloud to analyze the images using AI technology. We have trained Google Cloud to recognize any picture of rust corrosion and when spotted, the system alerts a Naval analyst to review and schedule the ship for repair. By leveraging Google Cloud's native computer-vision capabilities, the team successfully identified "corrosion of interest" in aerial images of vessels, with confidence scores of more than 90% and with very few false positives. This was an engineering feat that required complex integration between emerging software and hardware technologies, and has saved **the Navy thousands of hours a year in readiness**.

There are other examples of our work with the Department I would be happy to share – including using AI imaging to detect cancer, using AI to assist building simulation technology to train Air Force pilots and more.

As I mentioned, a critical use of AI is in cybersecurity solutions. While it is often hard to predict new kinds of attacks and new threats as they are constantly emerging, Google runs one of the largest and most secure networks in the world. Due to its scale and the threats it faces on a daily basis, we have a level of insight and visibility into the world of cyber threats, through all our global platforms, that allows us to assess and develop cutting edge defenses to whole classes of threats, not just particular attacks.

At Google Cloud, we have leveraged this expertise to deliver a new, unified AI experience through our Cloud services which give every data scientist, data analyst, and machine learning (ML) engineer the same tools we use at Google to secure their own networks. Like the Department of Defense, we must be constantly vigilant and ensure Google Cloud's security solutions and updates are informed by vulnerability and threat information as it evolves in real time. Indeed, as we have seen in many recent cyber attacks, some of the most dangerous attacks are those where multiple systems communicate in unforeseen ways to create chaos and wreak havoc. With this in mind, I'd like to offer the following observations and recommendations for how this committee can further support the Department of Defense in its mission using AI capabilities to secure its networks, applications and personnel:

1. **Using AI to defend against attacks**. As we have learned through recent events, our customers in the public and private sector increasingly understand that they must protect different parts of their network with different applications. There are known threat factors but all organizations must be able to spot new threat vectors that are constantly emerging and recognize that insider threats continue to be a real concern. DoD must stay on top of ensuring they have the right resources. And I'll attempt to illustrate this with how Google thinks about each of these threats:

    a. First, AI allows for monitoring known threats at a massive scale**.**

        i. Threat hunting and investigation tools are used to look at historical data and determine if exploitation was attempted - or they can be used as vehicles for monitoring active exploitation.

        ii. On-demand scanning of containers (containers are isolated software packages that contain everything the software needs to run).

    iii. Active scanning that detects Domain Name System (DNS) calls to known malicious sites (the DNS is effectively the "phonebook" of the internet).
    iv. Tools to detect common exploit attempts.

b. Second, AI excels at anomaly detection and emerging threats.
    i. Implementing passive detection rules in Event Threat Detection (ETD) and Security Health Analytics.
    ii. Tools to detect potential attacks include using custom reports in Edge API Analytics. (API stands for Application Programming Interface, which is a software intermediary that allows applications to talk to one another)
    iii. Tools to create web application firewalls as layered defenses to protect against attacks until all vulnerabilities can be patched.

c. Finally, AI can assist in identifying insider threats. AI is particularly best suited to identify insider threats because it has the capacity to analyze billions of parameters an hour. The need to protect against insider threats is also part of the Administration's push toward agencies embracing a Zero Trust philosophy.

It is worth noting that AI is trained and powered by data and so having accurate, well curated sources of data is key to threat hunting. For example, tools like VirusTotal provide threat context and reputation data to help analyze suspicious files. These tools use live flux samples of data against historical data in order to track evolution of certain threat actors, malware families and automatically generate "indicators of compromise" to protect organizations.

2. **Breaking down data silos to harness the full power of AI.**
Today, data exists in many formats, is provided in real-time streams, and stretches across many different data centers and clouds all over the world. From analytics, to data engineering, to AI/ML, to data-driven applications, the ways in which we leverage and share data continues to expand. Data has moved beyond the data analyst and now impacts every employee, every customer, and every partner. With the dramatic growth in the amount and types of data, workloads, and users, we are at a tipping point where traditional data architectures — even when deployed in the cloud — are unable to unlock their full potential. As a result, the data-to-value gap is growing.

Insights are not just locked in raw data — they're locked in data from many sources and silos — meaning the ability to unify datasets is a prerequisite to applying AI, in a structured and purpose-built manner, to applications. There are many opportunities to ensure the Department can operate different services across different and disparate data networks. For example, Joint All Domain Command and Control (JADC2) is seeking to do just this by allowing information sharing through interfaces and services across all domains. AI can enhance the security of this effort and ensure that the Department is reviewing the data for learnings, anomalies, changes and patterns.

A great example of this is how we are using AI systems for anti-money laundering and countering the financing of international terrorism ("AML/CFT"). Money laundering fuels drug trafficking, human trafficking, and terrorist activities. AI-enabled AML/CFT approaches, on the other hand, can develop a much more sophisticated analytic lens capable of ingesting massive volumes of data, in a more timely way, to detect new patterns and anomalies that might bypass simple, rules-based logic. These engines can be trained to improve accuracy, reduce false-positives, and help perform internal risk assessments and better determine when, amongst millions of legitimate transactions being processed, bad actors are trying to move criminal money. AI can further incorporate more contextual signals and generate more targeted flags for investigators, reducing toil and allowing them to focus on the most serious issues that are identified. AML highlights the opportunities this committee, the Department, and the private sector can focus on as we ensure the United States continues to lead in the development and deployment of artificial intelligence.

At Google Cloud, we have made it a priority to deliver cutting-edge cloud-native capabilities for distributed workloads spanning public cloud, private cloud, and multi cloud environments. Additionally, managing data across disparate locations creates silos and increases both risk and cost — especially when data needs to be moved. Innovations such as data lakes offer the ability to unify data stored across multiple cloud providers without worrying about the underlying storage format or system, which eliminates the need to duplicate or move data, which in turn reduces cost, inefficiencies, and security risks. This approach permits innovation by using multiple vendors, clouds and technologies, but it also increases competition and will likely lower prices for the Department and taxpayers.

But, it is not just about ensuring we have thousands of databases and data tables. The personnel at DoD must have the proper skills and training to capitalize on these insights. If an AI system identifies 27 new threats, we need DoD teams sitting inside the Department to quickly prioritize and address the threats. This is a vastly different way of thinking than the traditional "waterfall approach" which involves slower, deliberate planning and can constrain the more agile type work that is necessary in these scenarios. This is a classic challenge in large bureaucratic organizations. At Google, by the time a threat is discovered, we need to have a patch in place well within 24 hours. In two weeks, we need to have developed a permanent solution, and shortly thereafter, we need to have a post-mortem which describes the event and includes a recap of the timeline, description of user impact, root cause, action items and lessons learned.

3. **Capitalizing on data insights through human-machine teaming.**
   To understand the full opportunities of AI in DoD's mission, it must also ensure the Department can inject AI into its workflows. Understanding of AI-based tools cannot be limited to those with programming skills only.

   To be clear, this is not a procurement issue. Instead, what is needed is leaders to think about whether AI tools within the Department can help solve the challenge. Usually the answer is yes. Then the Department must have the ability for teams to quickly build/adapt/leverage an AI system — in hours or days — to address problems like finding a ship lost at sea or responding to an active threat event. Vertex AI and AI infrastructure provide tools for data scientists to build custom AI for their own problems at scale. Today, AI platforms like ours require nearly 80% fewer lines of code to train a model with custom libraries and data scientists can now build and train models 5X faster on Vertex AI than on traditional notebooks.

   Human teams, such as those formed by analysts and data scientists, must have a common understanding and opportunity to bring machine capabilities into the mission by building out an end-to-end AI experience where they can extract value from data and use AI out of the box to maximize value at a moment's notice. Imagine for a moment that there were different types of databases across the department that track shipping container movements around the world. Then imagine that another database holds information about the contents of each container and yet another that can analyze components or materials used in individual products inside the containers. Brought together, an AI system then identifies that there is a particular metal alloy used in each of the products that all appear to be heading to the same country in different ports. Cross-linking and joining data in this manner allows for constant pattern detection for unexpected defensive concerns and can help analysts identify emerging trends from data across different departments in new and novel ways.

This is especially important as our adversaries will continue to look for gaps in systems – including AI systems – that may be exploited in both simple and complex ways. The term "adversarial AI" may be known to you already but it is an increasing area of research. As I mentioned earlier, the most dangerous attacks are those where multiple systems communicate in unforeseen ways to create chaos and wreak havoc. AI is further enabling these kinds of attacks, but it can also help defend against them.

In the last several years, researchers at Carnegie Mellon proved that AI can act in super-human ways. This was recently demonstrated in a straightforward game of poker. Operating on incomplete information and against multiple parties, the system beat leading professionals by bluffing and misleading human adversaries. This is an indication of more to come. The poker demonstration offers valuable insights into the future of cyberdefense and warfare: our adversaries will continue to understand new and novel ways to leverage AI to mislead and attack.

As you can see, from poker games, to thwarting money laundering, to protecting networks from cyber attacks, to spotting corrosion on the hull of Navy ships, AI can be used to spot patterns and anomalies generally faster and with more precision than humans. AI technology can help the Department scale its analysis of these patterns and anomalies for threats and learnings. I urge the Department to embrace AI, particularly in its efforts to secure its networks.

Let me conclude by recognizing the importance of the work of this subcommittee, and its efforts to ensure the United States remains a leader in AI and cybersecurity, given the increasingly complex landscape. With AI, the work of 5,000 people can become the equivalent of 50,000.

My hope is that the Department will continue to make the right investments in training, technology, and management that will facilitate more experimentation, prototyping, and execution that will be necessary. It is also critical that the Department continues to make comprehensive technology investments - in cloud migration, data set curation, API management, network connectivity to increase operational effectiveness and deliver proven innovation.

We all have a role to play to prevent and detect threats online. Being transparent with governments, customers, and government entities when it comes to cyberattacks is one of our key principles and is critically important when responding to incidents at scale. I suggest this committee continue to encourage the use of modern, cloud-based technologies to improve long-term security, based on investments in defense-in-depth. Diversity in the ecosystem, especially with cloud-based solutions, reduces overall risk and fosters and improves resilience against attacks. In addition, products and services that enable portability and interoperability foster resiliency.

Thank you for the opportunity to speak with you today. I look forward to continuing to work with Congress on these important issues, and I'm happy to answer any questions you might have.