OCTOBER 24, 2024

Background Press Call on the U.S. Approach to Harnessing the Power of AI for U.S. National Security

Via Teleconference

MODERATOR: Good afternoon, everyone. Thanks so much for joining today's call to discuss the U.S. approach to harnessing the power of AI for U.S. national security, ahead of tomorrow's release of the National Security Memorandum.

As a reminder of the ground rules of this call, this call is on background, attributable to senior administration officials, and it is embargoed until 6:00 a.m. Eastern on Thursday, October 24.

For your awareness, not for your reporting, on the call today we have [senior administration official] and [senior administration official].

Following the call, we'll provide you all with some materials under the same embargo, so be on the lookout for those.

Our speakers are going to have a few words at the top, and then we'll turn it over to some of your questions.

With that, [senior administration official], I'll turn it over to you.

SENIOR ADMINISTRATION OFFICIAL: Thanks, Eduardo. And thanks to all of you for joining us this evening.

So, we're really pleased to report that tomorrow we'll be releasing a National Security Memorandum on Artificial Intelligence signed by the President.

And we want to start off just by sharing a little bit of context for this, which really begins with the fact that the United States has a very strong hand in AI

today. We design the most advanced hardware. We host the leading AI companies that are building the most advanced AI systems, and really have a dominant market share in artificial intelligence globally. And thanks to the President's CHIPS Act, we are building more resilience in our chip supply chains as well.

But as many of you know, the innovation that's happened, particularly in this current wave of frontier artificial intelligence, has really been driven by the private sector. And it's critical that we continue to both foster that leadership but ensure that the government, and particularly with this National Security Memorandum, ensure that our national security agencies are adopting these technologies in ways that align with our values.

And a failure to do this, a failure to take advantage of this leadership and adopt this technology we worry could put us at risk of a strategic surprise by our rivals, such as China.

And as you all know, there are very clear national security applications of artificial intelligence, including in areas like cybersecurity and counter-intelligence, not to mention the broad array of logistics and other activities that support military operations.

Because countries like China recognize similar opportunities to modernize and revolutionize their own military and intelligence capabilities using artificial intelligence, it's particularly imperative that we accelerate our national security community's adoption and use of cutting-edge AI capabilities to maintain our competitive edge.

So, President Biden's first-ever executive order, signed last October, on artificial intelligence was a key step forward to ensure that America leads the way in seizing the promise and managing the risks of AI.

In that executive order, the President specifically directed the development of this National Security Memorandum to ensure that we maintain our edge over rivals seeking to leverage AI to the detriment of our national security, while also building effective safeguards to ensure that our use of AI upholds our values and preserves public trust.

So, consistent with the President's direction, we've been engaged in a policy process over the last year or so to advance those aims and complete this

National Security Memorandum.

And tomorrow, the National Security Advisor, Jake Sullivan, will deliver remarks to rising military and intelligence professionals at the National Defense University so he can speak directly to the very national security professionals and leaders who are going to be implementing the core of this strategy.

During his remarks, Jake will talk about what led us to this moment in artificial intelligence, both in terms of its development and our views on why it is so critical for national intelligence and why, therefore, the President has issued this National Security Memorandum on AI.

Jake will also outline how the United States must strengthen our own advantages in artificial intelligence, how to harness that advantage in a responsible manner for national security, and also how the United States can do this work in lockstep with our partners around the world in ways that will protect our national security while also leveraging our advantages in AI for the benefit of countries around the world.

So, we hope you'll join us for those remarks as well.

With that, I'll turn it over to my colleague to provide more detail about the NSM itself.

SENIOR ADMINISTRATION OFFICIAL: Great. Thanks. And thanks, everybody, for joining.

As many of you know, the administration's approach to AI is rooted in the premise that capabilities generated by the transformer and large language model revolution in AI, often called frontier AI, are poised to shape geopolitical, military, and intelligence competition.

Now, most of the NSM is unclassified and will be released publicly. It also contains a classified annex that primarily addresses adversary threats.

Now, the principles guiding our work in the NSM are simple. They are that the U.S. should first lead the world's development of safe, secure, and trustworthy AI, and establishing a stable and responsible framework to advance international AI governance. And as a result, the NSM serves as a

formal charter for the AI Safety Institute in the Department of Commerce, which we have created to be the primary port of call for U.S. AI developers. They have already issued guidance on safe, secure, and trustworthy AI development and have secured voluntary agreements with companies to test new AI systems before they are released to the public.

Second, another principle is that the U.S. should harness the most advanced AI systems with appropriate safeguards to achieve national security objectives. And we are directing that the agencies gain access to the most powerful AI systems and put them to use, which often involves substantial efforts on procurement.

And finally, all of this must be done in accordance with our values.

So, alongside the National Security Memorandum itself, we are publishing a companion document called the Framework for AI Governance and Risk Management for National Security that provides guidance on how agencies can and cannot use AI.

So, we also believe that we must out-compete our adversaries and mitigate the threats posed by adversary use of AI.

So, in summary, what I've outlined are essentially three core principles that you'll see throughout the documents: securing the U.S.'s lead on AI; two, harnessing AI for national security; and, crucially, building in the governance framework to ensure that we are actually accelerating adoption in a smart way, in a responsible way, by having clear rules of the road.

With that, I'll turn it over to Eduardo.

MODERATOR: Thank you both. We'll now turn to our Q&A portion. If you'd like to ask a question, please use the "raise your hand" feature on Zoom.

First up, we'll go to the line of Katrina Manson. You should be able to unmute yourself.

Q Hi there. Thanks so much. I would love to ask how you see the U.N. intention to have countries sign up to a ban on lethal autonomous weapons by 2026 and if any of your work foresees the U.S. signing up to that.

Many of the harms that you try to prevent on the civil use of AI, obviously in terms of bodily harms, are very much implied with the use of AI for the military. And in the case of Maven, AI targeting is already being used to support battlefield firing in the Middle East by the U.S. Can you address the very serious safety concerns around the use of AI targeting and whether you will consider a ban on lethal autonomous weapons, which can use AI?

SENIOR ADMINISTRATION OFFICIAL: Thanks for that question. I'm happy to start with that.

So, first point is, as I think [senior administration official] noted, we'll be releasing tomorrow, alongside the National Security Memorandum, a framework on responsible use of artificial intelligence in a national security context. And so, you'll see there really a lot of detail on kind of all the steps that we're taking to ensure these systems are used responsibly.

Now, and the other thing I would point out is: While it's not necessarily part of this NSM, although there's a nod to kind of our diplomatic efforts and kind of direction to double down on those, some of you may be aware of the Political Declaration on Responsible Military Use of Artificial Intelligence and Autonomy. And that's a declaration where the Vice President, in fact, has kind of taken a leadership role. And we have around 60 countries that have signed up to this declaration, which is really focused squarely on how AI and autonomy should be used. And most recently, there was a summit held on this by South Korea.

So that's another area where that combines both the substance that you'll see in the framework on responsible use, but also, really, diplomatic efforts that we've been leading over the last few years.

SENIOR ADMINISTRATION OFFICIAL: And, sorry, if I can add to what was just mentioned. The framework itself you'll see actually references the political declaration that was just mentioned, and it also outlines the requirement for adherence to the Department of Defense's Directive 3000.09 and successor related policies that address autonomous or semiautonomous weapons systems.

But in addition to that, as was just mentioned, there are a number of outlined prohibited use cases, as well as high-impact use cases that are relevant. And

one theme you'll see in both the NSM and the framework document is the fact that we need to ensure that AI is used in a manner consistent with the President's authority as Commander-in-Chief to decide when to order military operations in the nation's defense, for instance.

MODERATOR: Thank you. Next up, we'll go to the line of Garrett (inaudible). You should be able to mute yourself.

Q Hello. Can you all hear me?

MODERATOR: We can, yes.

Q Great. You mentioned that some of the commitments from companies are voluntary. And, you know, just covering the big fight around legislation here in California, companies seem, from my perspective at least, to very much want to keep those commitments to safety and that kind of thing voluntary, rather than sort of required or legislated.

And I'm just wondering if, you know, the administration has a view, or if it's published as part of this, about trying to sort of codify those voluntary commitments and make them more, you know, ironclad and not sort of up to the whims of these CEOs.

SENIOR ADMINISTRATION OFFICIAL: Thanks, Garrett. So, I think on that point, I would just say we continue to work with colleagues on the Hill. There are a number of proposals relating to, you know, regulations on artificial intelligence. And so, that's really — that's, really, ongoing.

I think, really, the emphasis in the National Security Memorandum is really kind of making commitments ourselves as a government about how we will adopt and use artificial intelligence. You know, as you point out, we have played a leadership role in getting some of those commitments from the companies. We have taken those commitments and kind of — to the international stage, through the G7 and the Hiroshima process as well.

But, really, what we're focused on tomorrow is what commitments can the government itself make on responsible use, which we think is important, by the way, not just for its own sake, but we also think that's important to enable us to both accelerate both the development and also accelerate the adoption

of use as well. And that's a point that I think you'll hear the National Security Advisor focus on as well tomorrow.

MODERATOR: Thank you. And next up, we'll go to the line of Patrick Tucker. You should be able to unmute yourself.

Q Hi. Thanks. Pat Tucker from Defense One.

There's a new paper out, actually this week, from Meredith Whittaker and a couple other folks at the AI Now Institute, actually pointing out some of the potential dangers of some of these commercially facing AI products in national security contexts.

And they point out that some of these generative AI tools have very large — unacceptably large false positive rates. They hallucinate, often, a lot. And sometimes to train them, they rely on publicly available data, including data that might come from data brokers and other sources that poses a potential privacy risk, particularly to Americans, because Americans produce a lot more purchasable data than do citizens in China or Russia.

So can you talk a little bit about how this memorandum does or does not address data vulnerability of Americans and some of the potential risks in the national security setting of adopting commercial and consumer-facing AI tools that have high hallucination rates or false positive rates? Thank you.

SENIOR ADMINISTRATION OFFICIAL: Do you want to start with that? You can join as well.

So, thanks for the question. Look, I think some of these, you know, concerns I think are ones that I think colleagues in the national security community are acutely aware of. You know, there are a few points here.

One is, you know, we have to go through a process of accrediting systems. And that's not just for AI systems, but you know, national security systems generally. And so, that's point one, to kind of ensure that they are fit for the purpose or particular mission.

I think the second point is: We are, you know, very — I think very aware that what we're doing at this stage is really trying to ensure that we have pilots

and some important experimentation happening, because there are going to be challenges associated with adopting any new technology.

Third is, the framework that [senior administration official] mentioned is one that's going to have to be continuously updated. And we have tried to set it up in a way so that that can happen in real time as there are challenges that are inevitably encountered.

And parallel to the policy process here, we have a lawyers group that is kind of working very intensively to ensure that, obviously, all existing law is complied with, but also to ensure that novel legal issues as we encounter them are addressed in a timely way as well.

I do want to just address the point on data that you mentioned specifically, which is, you know, we have been very concerned about the ways in which Americans' sensitive data can be sold, really through the front door — through first collected in bulk, then sold through data brokers, and then end up in the hands of our adversaries. And so, that's something that the President issued an executive order on to try to restrict adversary access to some of that data. And, in fact, just this week, we took one more step in the regulatory process through a notice of proposed rulemaking to try to get that final later this year.

SENIOR ADMINISTRATION OFFICIAL: And if I can just add on that.

So, in addition to the work that the AI Safety Institute is going to do, and as [senior administration official] mentioned some of the other work, you'll see that in the NSM itself there are very specific requirements for specific agencies and our intelligence community, and, for instance, the Department of Energy to do classified testing of different systems for different purposes for this very reason.

And in addition to that, as [senior administration official] mentioned, there's a strong focus on experimentation here for this very reason. We want to see rapid adoption, but we also want to see experimentation that will tease out kind of what missions are best suited for various systems and also tease out the challenges of them. And that's going to require leaning forward and experimenting, adopting, and then doing all of the work that was just mentioned as well, in terms of both policy and legal review.

MODERATOR: Thank you. We have time for one more question, and we'll go to the line of Maria Curry. You should be able to unmute yourself.

Q Hey. Thanks for taking my question. I'm wondering if export controls are part of this at all. And if so, can you elaborate how those might be helpful?

And then, if you could just elaborate, too, on the third point. Could you dig in a little bit deeper into how agencies can or can't use the technology? Could you provide an example or two of that? Thank you.

SENIOR ADMINISTRATION OFFICIAL: I can speak to the export control piece, and, [senior administration official], maybe you can speak to some of the prohibited use cases.

So, really, the NSM does kind of address, kind of as a matter of policy, the importance of protecting advanced AI technologies so that they're not used against us by adversary militaries or intelligence services. And so, at a high level, it does kind of try to emphasize the importance of maintaining those policies and making sure that we are continuously adapting to efforts to circumvent those measures.

And as you know, those export controls cover not only GPUs, the advanced AI chips, but also the semiconductor manufacturing equipment that's necessary to manufacture those as well. So, that full aspect of the supply chain.

[Senior administration official] do you want to say anything about prohibited uses?

SENIOR ADMINISTRATION OFFICIAL: Sure. So, you'll see in the accompanying framework document that I mentioned, it identifies both prohibited, as well as what we call high-impact AI use cases, based on the risk that they pose to national security, international norms, democratic values, human rights, civil rights, civil liberties, privacy, and safety.

And on the prohibited end of the spectrum, these will be — not surprising, but there are clear prohibitions on use of AI with intent or purpose, for instance, to unlawfully suppress or burden the right to free speech or the right to legal counsel.

There's also prohibited use cases around, for instance, removing a human in the loop for actions critical to informing and executing decisions by the President to initiate or terminate nuclear weapons employment, for example. That runs the spectrum of kind of military-related activities, but also protecting civil liberties and tracking international norms.

But in doing that, we actually view these restrictions — so these prohibitions, for example, as well as the high-impact cases — as being important in clarifying what the agencies can and cannot do. That will actually accelerate experimentation and adoption. Because one of the paradoxical outcomes we've seen is: With a lack of policy clarity and a lack of legal clarity about what can and cannot be done, we are likely to see less experimentation and less adoption than with a clear path for use, which is what the NSM and the framework tries to provide.

MODERATOR: Thank you. That's all the time we have for today. Big thanks to our speakers, and thanks to you all for joining.

As a reminder, this call is on background, attributable to senior administration officials. And this call and its contents are embargoed until 6:00 a.m. Eastern tomorrow.

Thanks, all, for joining. And be sure to tune in tomorrow to National Security Advisor Jake Sullivan's remarks on this topic. Thanks again.