wiley

AI: The Next Big Thing in Government Contracting

—

Wiley Government Contracts Podcast
June 22, 2023

Wiley Government Contracts · AI: The Next Big Thing in Government Contracting

On this episode of Wiley's Government Contracts podcast, Privacy, Cyber & Data Governance partners Duane Pozza and Kat Scott join Government Contracts partner Tracye Howard and host Craig Smith to discuss how federal policy on artificial intelligence (AI) is affecting government contracting. Join the conversation on the latest AI trends, government efforts to regulate AI development, and how contractors can mitigate the risks of future regulatory efforts.

Transcript

—

Craig Smith

Welcome back everyone to our Government Contracts podcast series. Today we want to talk about Artificial Intelligence – AI – it's everywhere you look it seems. You can't go a week or a day or an hour without a news article or a tweet about it. There's new and innovative uses of AI, and, of course, there are risks that come along with it. Not just the ones you see in movies, but the ones that counsel and all sorts of compliance personnel have to think about everyday. So how do you deal with that? How do you identify the risks and figure out how to manage and mitigate them? Today I have three

Related Professionals

—

Tracye Winfrey Howard
Partner
202.719.7452
twhoward@wiley.law

Duane C. Pozza
Partner
202.719.4533
dpozza@wiley.law

Kathleen E. Scott
Partner
202.719.7577
kscott@wiley.law

Craig Smith
Partner
202.719.7297
csmith@wiley.law

Practice Areas

—

Artificial Intelligence (AI)
Government Contracts
Privacy, Cyber & Data Governance

guests – very lucky today – my partners Tracye Howard, Duane Pozza, and Kat Scott – to talk about trends with AI policy, and we're going to focus in on, both broadly or focus broadly and then zoom in on government contracting issues. When I think about AI, one thing that occurred to me in getting ready for this podcast is there are some themes that sound reminiscent from prior, it's the next big thing in federal contracting like, data analytics, but people don't make movies about SQL servers and Oracle databases. So AI has, I think, a broader kind of attention to it that is really important to understand before thinking about how is this really going to affect government contracting, understanding the broader principles are really important, and so for that, Duane, I want to get started with you. What's going on with AI at the policy level in the federal space?

Duane Pozza

Thanks, Craig, and happy to be on the podcast. The short answer is there's a lot going on at the federal level in the AI policy space. I think it might be interesting to start sort of at the top with the White House. The Biden Administration has continued what I think is the previous administration's push to really sustain, enhance, and promote AI research and development to make sure the U.S. remains a leader in AI development. It has done this through a number of different work streams, lots of different acronyms coming out of the White House and various AI task forces to try and promote this kind of development. But at the same time, there has also been a focus on the potential risks, and I think a very prominent discussion about what federal AI policy should be in terms of managing the risks, at the same time promoting the potential benefits of the AI. So, just to take a few examples, you've seen some high-level meetings at the White House or convened by the White House with leaders, developers of AI to talk about different kinds of you know voluntary measures they could take potentially or even potential legislation to deal with AI on a large level.

You also see some more specific things coming out of different offices. So, for example, the White House Office of Science and Technology Policy has been focusing a lot on different AI work streams and recently released what they called an AI Bill of Rights. So that is a sort of consensus-driven document that lays out a bunch of different concerns that stakeholders might have when implementing AI, which is designed to allow it to be beneficial when it's used, but also to address specific things like transparency, privacy, and avoiding bias that are all part of implementing AI in a responsible and successful way. There's a few other work streams going on, and I'll turn it over to Kat actually to talk a little bit more about that.

Kat Scott

Yeah, thanks Duane. I agree, there is a lot going on at the federal level from federal agencies in the AI space. As folks know, last year the FTC took steps to launch a privacy and data security rulemaking proceeding. They released an ANPRM – Advanced Notice of Proposed Rulemaking – and it was really broad, it asked a lot of questions about commercial data practices, but one big area of focus and a clear area of concern from the FTC is AI or automated decision making, so to the extent that that rulemaking proceeding moves forward, I think we can expect continued scrutiny and potential rules for AI coming out of the FTC. NTIA is also an agency that has been really focused on AI recently. They're an agency within the Department of Commerce, and they have recently launched two proceedings to draft reports that may implicate and inform future federal privacy and AI policy. So, the first is a report on the intersection of privacy, equity, and civil rights

issues. They're asking questions about whether and how commercial data practices, including AI practices, can lead to disparate impacts and outcomes for marginalized or disadvantaged communities. And then the second is a report more focused on AI. It's AI accountability mechanisms. So, a lot going on from federal agencies taking a more regulatory approach to AI. And then on the flip side of that, you have NIST, another agency, and the Department of Commerce, and they recently finalized and released an AI risk management framework. So, this is a document that helps organizations identify and manage AI risks. And we think it's really exciting, right, it's taking less of a regulatory approach to AI and instead looking to facilitate innovative deployment of trustworthy AI, at the same time as identifying and mitigating risks.

Craig Smith

It sounds like a lot of agencies and acronyms that many of our listeners might not have as kind of their day-to-day agencies that they're focused on, and what I hear from you, and you can nod along or even say yes if you want, is there's going to be some push-pull that, Tracye, we've seen our contracting clients deal with where there's, you want consistency across government, but that takes a long time, and you can start to slide behind industry, and even in the defense sector, our adversaries. Seems like that's going to be a tension we're going to be facing from here on in.

Duane Pozza

Yeah, I agree with that, and you know one thing that OSTP – another one of those acronyms – is doing right now is they actually have a Request for Comment on development of a potential national AI strategy with, I think, one goal being trying to find a more uniform approach, across at least the federal government, in treating how AI will be affected.

Kat Scott

And we see that a lot in other spaces too – privacy, cybersecurity – this risk of fragmentation. Everybody's interested. Everybody sees the risks and the potential benefits, and everybody wants to get in on the action. Just adding to what Duane was saying, I think the other thing that we have to think about in this space is state action. States are not ignoring the AI trends either, and we're seeing state agencies and state legislatures looking to address the potential risks of AI too, so that all adds to this issue of fragmentation.

Craig Smith

That's a great point, Kat, and reminds me of the conversation we had a few months ago about state privacy laws.

Kat Scott

Exactly.

Craig Smith

It can get dizzying quickly trying to manage all of this. But then going from policy to what's showing up in my contract – so I'm a federal contractor, I might think about, well AI is going to affect me in one or more of three ways. It could be what I'm actually signing up on the contract to deliver. It could be something I use as a tool to deliver something else, like I use an AI program to engineer a change to a jet wing, or it could be what the government uses to either award or manage my contract, manage pricing. Those seem to me like the three basic ways that I might categorize the impacts on me, but when I sit down with my paper copy of the FAR and the DFARS, Tracye, what am I going to find there?

Tracye Howard

You're going to find nothing. And so, just to be clear, the government is already buying AI today, and you know DOD tried to do an inventory a couple of years ago, they came up with 700 or so projects that involve AI to some extent or another, and obviously other agencies have that too, and so the government's already buying this and, Craig, as you said, there's nothing in the FAR – in the Federal Acquisition Regulation, which is our bible of government contracts. There's also nothing in the DFARS, the DOD supplement to the FAR, and so we have all of these agencies doing different things and even not just across different agencies, but maybe within an agency, different buying offices are doing different things, applying different standards, having different requirements, and so I think it is a bit of an opportunity, some of these government-wide efforts that Duane and Kat were talking about to maybe bring some standardization to the process and have contractors not being forced to try to comply with different requirements across different agencies of the government. And so, what we've seen you know, this is similar to what we've seen in cybersecurity and some other spaces up to now, is contractors tend to be, in federal procurement generally tend to be sort of, I don't want to say the guinea pig, but kind of the first area that gets regulated or has standards applied to it, and this is because historically, often different administrations of both parties and Congress of both parties have trouble agreeing on things and how to apply them to sort of private industry writ large, and so they try to maybe apply some standards and requirements to government contractors as a starting place, with the hope that maybe that will trickle out into the broader economy, and so we've seen this for many years in the space of sort of labor and employment and wage and hour requirements. More recently we've seen it with cybersecurity, both with sort of breach reporting and then broader cybersecurity requirements, and now most recently we're seeing it with software development requirements, where the government is rolling out these, to some extent, aspirational requirements that they think everyone should comply with. They're requiring government agencies and contractors to comply with them and then putting them out as standards that, you know, they think are a good policy for others in the industry to comply with on a voluntary basis.

And so, bringing that down to a little more granular level, so OMB is going to be releasing a draft policy this summer on AI use across the government. Obviously, that's going to affect contractors. And so, it could affect contractors in sort of how they, as Craig was referencing, solutions that they're delivering to the government, and those standards that apply to the government would apply I think equally to whatever is being delivered by the contractor. It may also apply to those things that you're sort of using in the background, the second category that Craig was talking about, where you know these AI policies that apply to government agencies

might also apply to those things that contractors are using to develop the solutions for the government but not actually delivering to the government.

Craig Smith

And so all this, I mean I'm looking at all three of you because we are actually in-person. And AI can feel a little overwhelming, you know, you're responsible. It's not like cybersecurity we give to the infosec people, or an equal employment issue that's developing that goes maybe to HR or benefits or whatever. This can feel all encompassing. We talked about privacy and the employment aspects and, kind of, the FAR and DFARS don't really have a framework for say, a contracting officer to interact with. And I can think about all sorts of issues, like how do I explain decision making by an AI feature to an Inspector General agent who has questions about how we did something on a contract. But you got to start somewhere. And so, I'm a federal contractor, if I'm sitting here going, I'm not Joe or Jane Artificial Intelligence. What can I do today to start getting the, you know, information I need and start making sure my company is moving in the right direction from a policy perspective or a practice perspective to be caught up in where we need to be for AI looking ahead?

Tracye Howard

So, I think one thing to do is just assess what it is that you're using in your company today. Are there any products and services that you're either providing to the government or using for your own business purposes that include AI? Understand what it is that you have, where you're utilizing them, how you're utilizing them. And so, at least you'll know when these standards come out, if they come out, what they're going to be applied to, and then start to understand you know what it is that the government is buying today and what those standards are, and some agencies, the Department of Defense in particular, has put out, they had this online marketplace called Tradewind Solutions. They post their problem sets for AI and you know you can understand, here's what the government's currently seeking, and sort of where they're going and understand how that fits with your business.

Duane Pozza

I think another thing that companies can do is to really get familiar with and work to adopt and conform their existing practices with the set of best practices that I think you can find in AI already. I talked earlier about the AI Bill of Rights that came out of the White House, and that's sort of one roadmap. But an even more detailed roadmap is one that Kat mentioned, the AI Risk Management Framework that was developed by NIST. It has a lot of detail actually about the different ways that companies can actually take these high-level principles and boil them down into action items to actually work through some of these issues. And, again, you know at a high level, that includes things like transparency, explainability, privacy, security, ways to identify and avoid bias, and also accountability, which is the idea that there should be human oversight over AI in certain circumstances. And the last thing I'll say on that, and it's obviously a complex topic, is it is risk-based, so like a key aspect of this is identifying where the risks are highest in using AI, and where they might be lower, and then adjusting your practices accordingly, and I think also really thinking about the governance structure that goes around it. So, Craig, you talked about, you know, who is the person who's going to be responsible. Well,

I don't think there's one single framework, but ultimately there has to be a structure of folks who have insight into what's going on with AI and it probably needs to be a group of people who include technologists, probably some lawyers, and business folks who can all sort of look at what's going on and make sure that there's proper oversight as it's being rolled out.

Kat Scott

Yeah, I agree with all of that. I think the only other thing I would chime in here is that companies really should be looking inward, right? Even if you're not contracting with the government for goods or services that use AI, you should definitely be looking at your internal policies and procedures and making sure that they're set up to deal with internal uses of AI tools. So, for example, we recommend that organizations establish or update their policies and procedures to account for employee use of generative AI tools. This could include updating your organization's acceptable use policy or developing a standalone generative AI policy. We also encourage organizations to look at the representations they make to their customers to understand if those need to be updated to be transparent about any internal AI uses.

Craig Smith

I think that's a great point Kat, too, because it's almost timeless advice, because policies at least show hey we thought about this, and we're not just sending people off onto the internet with AI tools without any type of guardrails for them, even if they're not the perfect, at least starting somewhere is the place to go. So, Tracye, Duane, Kat, really appreciate, I think this is a conversation that will continue to evolve, and the good news is now that we have your voice prints, we can just generate a six podcast series for the rest of the year while you all are doing other things. So, I want to thank all three of you for joining us, really appreciate the insights, and thanks so much.

Kat Scott

Thanks, Craig.