***The Consumer Finance Podcast*: Navigating the NYDFS' Cybersecurity Guidance on AI**
Host: Chris Willis
Guest: Kim Phan
Recorded: November 4, 2024
Date Aired: November 14, 2024

Chris Willis:

Welcome to *The Consumer Finance Podcast*. I'm Chris Willis, the co-leader of Troutman Pepper's Consumer Financial Services Regulatory Practice. And today we're going to be talking about the New York Department of Financial Service's new cybersecurity guidance related to artificial intelligence, which is a matter of great interest to a lot of people in the industry.

But before we jump into that topic let me remind you to visit and subscribe to our blogs, TroutmanPepperFinancialServices.com and ConsumerFinancialServicesLawMonitor.com. And don't forget about our other podcasts. We have lots of them. We have the *FCRA Focus*, all about credit reporting. *The Crypto Exchange*, about everything crypto. We have *Unauthorized Access*, which is our privacy and data security podcast. *Payments Pros*, which is all about the payments industry. And our newest podcast, *Moving the Metal*, which is all about auto finance. Those are available on all popular podcast platforms.

And speaking of those platforms, if you like this podcast, let us know. Leave us a review on your podcast platform of choice and let us know how we're doing. And if you enjoy reading our blogs and listening to our podcasts, our mobile app is a great way to do it. It's available for both iOS and Android. Just head to your app store and type in Troutman Pepper, download it, and give it a try.

Now as I said today, we're going to be talking about some new guidance that came out on October the 16th from the New York Department of Financial Services dealing with cybersecurity and threats related to artificial intelligence relating to cybersecurity. And for a long time, the New York Department of Financial Services has really been at the forefront of cybersecurity, standard setting and regulation through its various actions. And they're remaining very much on the leading edge with this new pronouncement.

And today to discuss that with me is my partner and one of the co-hosts of our *FCRA Focus* podcast, Kim Phan. She's a very frequent participant on this podcast. Kim, welcome back. And thanks for being here today.

Kim Phan:

It's always my pleasure to be here, especially because there's so much going on in this space. I'm glad to come and tell your listeners about it.

Chris Willis:

Yeah. I think I'm really looking forward to this conversation. As I said, the New York DFS has been very proactive with regard to cybersecurity both in terms of the guidance and the regulations that it's issued, but also in terms of applying those in supervisory examinations of the entities that it has jurisdiction over. And now it's come out with this new guidance on October the 16th about AI and cybersecurity. Kim, tell the audience a little bit about what's in this new guidance. What obligations does it impose on the regulated entities subject to DFS's jurisdiction?

Kim Phan:

Well, to be clear to everyone who's listening, the guidance is just that guidance. It's not new rules. It's not new legal obligations. It's actually characterized as not creating any new obligations. But rather, it's positioning itself as providing insight. Merely letting financial companies know how New York DFS, if they were in their shoes, would be complying with some of their existing cybersecurity obligations when it comes to AI.

And the guidance leverages the risk-based approach that New York DFS has laid out in its cybersecurity regulation by making clear that New York DFS expects that financial companies, when looking at their cyber security maturity, are thinking about how they understand, assess and mitigate the specific cyber risks that are posed by AI which are substantial.

And regardless of whether or not a company plans to use AI for its own internal purposes, I think it's made pretty clear by New York DFS that they can't just ignore AI. Because even if AI tools are available to companies, and they're worried about how to deploy them, the bad guys have no such worries. They are definitely deploying AI. And those deployments by the malicious actors pose a very specific and dire threat to financial companies that aren't thinking about how to deploy countermeasures.

Chris Willis:

Yeah. I bet the bad guys' sort of modeled governance processes are not nearly as lengthier as involved as those in the financial services industry.

Kim Phan:

I would agree with you on that.

Chris Willis:

Even though it's couched as guidance, it's still something that I think the industry really still needs to take very seriously. Does the guidance identify any specific AI related risks that financial services companies need to be watching out for?

Kim Phan:

It actually looks at it from two different perspectives. It actually looks at AI risks posed by external threat actors, but also looks at the cyber risks caused by a company's own use or reliance on AI. And not everyone is doing that yet. But that is still something of concern to New York DFS.

Tackling the external risks first. The New York DFS identifies two specific risks. AI-enabled social engineering and AI-enhanced cybersecurity attacks. Let's talk about social engineering first. This is probably something you've all heard about a ton in the news already. AI being used to create highly-personalized mimicked voices or images. Deep fakes.

When, in prior days, before AI, that was just simple phishing. And then they had telephone deep fakes where it's grandma calling up and saying, "Hey, my lights are going to get cut off. You've got to send money right away." Or granddaughter calling and say, "Hey, I'm in a Mexican prison. You've got to wire money right away or I'm going to be here overnight." That's characterized as vishing.

We're now seeing text deep fakes. SMishing. And though this isn't addressed in the guidance, we're seeing AI being used to generate fake QR codes. The ones you scan in stores and other places. It's called qishing. And all this is being used to lure individuals to devolve sensitive information or to process fraudulent transactions. Wire money to fake accounts during a mortgage closing. Something like that.

But the New York DFS thinks about AI also more broadly rather than specific strategies very much aware that AI can be used to enhance the efficacy of cyber attacks because of the potency, scale and speed that AI can be used to analyze vast amounts of data. It can help identify and exploit security vulnerabilities in a company systems more quickly than a company can patch them. It can help develop malware variants and develop ransomware tools to bypass security systems.

And once they get into a system, the AI can be used to essentially conduct reconnaissance. Just wander around and just analyze every piece of data it can find so that it can give the bad guys better strategies to know when to strike. It could be in the system for months and it doesn't strike until the days right before, say, a large transaction or deal closing. Something like that.

From an internal perspective, the New York DFS is very concerned about supply chain vulnerabilities. The reality that more and more companies may become AI-dependent and cause potential business failures. Whether or not their own deployment of AI or their use of third-party service providers who may be using AI on their behalf. Each of those entities in the New York DFS's eyes is a potential link in the supply chain that could be creating AI vulnerabilities to cybersecurity risk. New York DFS expects companies to be looking at both external and internal risks when thinking about how to approach AI from a cybersecurity perspective.

Chris Willis:

Kim, with regard to those external risks and the sort of AI-enhanced confidence schemes so to speak, was the New York DFS guidance focused on those as being directed at, say, a financial institution's employees who might then give access to a system with PII on it, et cetera, to threat actors? Or did they also discuss the possibility that a financial institutions customers might be the targets or victims of those kinds of advanced phishing and similar schemes?

Kim Phan:

The focus is primarily on companies themselves and their ability to implement internal defenses with regard to warding off those types of schemes against their own employees and against their own systems. Now I'm sure the New York DFS is very interested in consumer education and ensuring that consumers themselves are protecting themselves against those types of risks. But that's not the focus of the guidance.

Chris Willis:

Okay. Got it.

Kim Phan:

Now that this guidance has been released, and it's a friendly reminder from the New York Department of Financial Services, what's your view on what financial institutions subject to DFS's jurisdiction should do in response?

Kim Phan:

Friendly at this stage, but may not stay friendly for long. The reality is some of the recommendations that New York DFS offers are, let's say, some more practical than others. And let me give you a couple examples. The basis for everything that New York DFS expects is a risk assessment. Now this is something that companies are expected to do. It's a building block for their entire cybersecurity programs. And it's supposed to evaluate different risks including those posed by AI and to match those risks against defensive measures that they would implement.

Assuming that you're doing your risk assessments, which is, again, the building block for everything, some of the more practical advice that New York DFS offers is with regard to those third-party service providers that we were just talking about who may be utilizing AI to help financial institutions operate. One of the more practical things they recommended was making sure that there are additional AI representations and warranties being added into those vendor contracts not only to understand what AI is being used on a company's behalf, but also what protections are in place for those organizations should a vendor's use of AI, say, pose fair lending risk or some of the other risks not just in the cybersecurity area.

Some of the advice that I think is a little less practical, one of those areas is with regard to access controls. New York DFS has been saying for some time now that multi-factor

authentication is one of the best controls you can implement as far as protecting your system. But here, they specifically said multi-factor authentication, for your audience that doesn't understand that, is something you know which is a password. Something that you have. A passcode that's sent to your phone that you use to log in. Or something that you are. Biometrics essentially.

And with regard to those two second factors, something that you have and something that you are, I don't know the New York DFS is coming up with solutions that are practical for implementation. They specifically said that AI can be used to deep fake things like SMS text messages, voice, video and other types of digital-based certifications and physical security keys so that companies should avoid using those tools when setting up their multi-factor authentication. And with regard to biometrics, they specifically said that companies should be thinking about not only a single-factor biometric. But to have biometric modality. Multiple layers of biometrics at once. Fingerprint in combination with iris recognition. Or fingerprint in combination with user keystrokes. Or if having only a single biometric authenticator, using deploying technology that detects liveness and texture analysis. It's not just the picture of someone's iris. It detects whether or not it's an actual live person that there's blood running underneath the iris. And they're confirming that this is a live person and not a fake.

How available those technologies are and how easy they are to deploy? I don't know that that is ubiquitous enough at the market now to say, "Hey, don't send codes to people's texts and SMS phones anymore or to their email addresses. Because that's too easy for bad guys to get around." I don't know that there's an easy solution around those that New York DFS is contemplating here.

Chris Willis:

Well, frankly, Kim, it doesn't seem like it would even be that easy if you had – let's say you had some mechanism to differentiate a picture versus a live image. I mean, what's to stop an AI from being able to replicate that live image? That seems like it would be something that would be much faster to develop than the countermeasures to it. And it just sort of results in this arms race where the bad guys are perpetually trying to keep ahead.

Kim Phan:

Well, that's always been the case, that the bad guys have better technology than the good guys. That's just the reality of the cybersecurity risk. At this stage, it's mostly whack-a-mole, right? Where you're trying to defeat things as you learn about them. That's some of the importance around information sharing amongst the industry, which New York DFS has long advocated that there be more sharing. So that if a couple of mortgage companies get hit at once, that they can share that information and maybe the next mortgage company is able to deploy a defensive measure. Or if depository accounts are being hit with a similar type of strategy by the bad guy, if they share that information, the next depository account can be defended by that particular bank.

That's something to be thinking about. But, honestly, that hasn't worked great in the past. Companies don't want to share when they've had problems. And they don't want to give away

basically the keys to their kingdom to their competitors to let them know how they're staying ahead of some of these threats.

But one of the things that I thought that was a better recommendation by New York DFS was the idea of implementing and establishing proactive measures to test for AI risks. Now they didn't use this term specifically. But the idea of running tabletop exercises with an AI focus. Tabletop exercises, Chris, again, for your audience who's not familiar with that term, are basically simulated data breaches. Where you help monitor and walk a client, one of our companies, financial institution companies, through what would happen if there were an AI-based threat to their organization. And they would talk through how they would think about addressing that. And that's one thing that a company can do right away to start thinking about how to address AI risks.

Another really good recommendation I thought by New York DFS and I think is more technically feasible was the reality that a lot of employees, whether permitted or not, are probably using ChatGPT. Whether or not it's on their actual work devices or just on their own personal devices, on their phone, New York DFS had suggested that financial institutions monitor for any unusual query behavior that are being submitted through ChatGPT to identify when someone's trying to extract large amounts of non-public personal information or blocking personnel that might be using AI in a way that could be exposing MPI into the public sphere. That seems like a pretty prudent and easily deployed technology.

Chris Willis:

The department has set forth a list of potential steps for companies to take about this issue. How confident can we be that if a financial institution takes all of these steps that that'll be sufficient for New York DFS in terms of their next examination or next encounter with the department?

Kim Phan:

That's a tough one, Chris. The reality is that what is reasonable today may not be reasonable tomorrow it may not be reasonable in the eyes of New York DFS ever. They are constantly seeing that the line in the sand is moving. It's very much a shifting expectation with regard to the types of defenses companies have to have a place. And the guidance actually makes clear. It's highlighting some of the threats that it has identified by its own cybersecurity experts. But it's not intended to be an exhaustive list of the types of threats companies might face and the types of defenses they might want to put in place.

You want to think about it the same way you would think about any other potential risk for a financial institution. You want to be thinking about who are the internal stakeholders that need to be pulled in to discuss this? New York DFS states specifically that its expectation is that senior leadership have some visibility into some of these AI risks. Companies need to be thinking about their own internal use cases as well as what external uses they make of AI.

If you're deploying an AI chatbot on your website to engage with customers, how vulnerable is that chatbot to interception or manipulation by a bad actor who might be capturing some of the

information a customer is telling you when they're communicating with you online? And thinking about what AI vendors are using on behalf of the organization.

The New York DFS, I mentioned earlier, that your risk assessment is really the building block for everything you should be doing. Every time you deploy as an organization, new AI. Anytime you're onboarding a vendor that is offering an AI tool. The New York DFS expectation is that you conduct a new risk assessment of that new tool, that new vendor. And be thinking about how that impacts your cybersecurity.

Chris Willis:

Kim, in closing, after having gone over some of these details about the New York DFS guidance, is there any sort of further words or parting thought you'd like to leave the audience with?

Kim Phan:

Yeah. Chris, I just wanted to add. We've talked a lot today about the defensive posture that companies have to take with regard to AI. But I want to make sure that we're also making clear that there's all kinds of substantial cybersecurity benefits that can also be gained by integrating AI into a financial institution's cybersecurity tools, controls and strategies. The same way the bad guys can use AI to analyze fast amounts of data, so can the good guys. To automate routine, repetitive tasks such as reviewing security logs and analyzing behavior, detecting anomalies. Trying to efficiently patch vulnerabilities and threats. Responding quickly once a threat is detected. And expediting the recovery of normal operations by utilizing some AI tools. I want to leave the audience on a bright point and give them something to look forward to that it's not all bad news. There are many, many good things that can come from AI.

Chris Willis:

Thanks, Kim. I think that's great insight and advice. And, obviously, here at *The Consumer Finance Podcast* and here at Troutman Pepper, we're going to continue to very closely follow regulatory pronouncements and other developments with respect to AI and cyber security. And, of course, bring those updates to our clients and friends in the industry as they occur.

Thank you, Kim, for being on the podcast today. And, of course, thanks to our audience for listening in to today's episode as well. Don't forget to visit and subscribe to our blogs, TroutmanPepperFinancialServices.com and ConsumerFinancialServicesLawMonitor.com. And while you're at it, why not head over and visit us on the web at troutman.com and add your email address to our CFS mailing list? That way, we can send you copies of our alerts and advisories and invitations to our industry-only webinars that we do from time to time.

And as I mentioned at the top of the show, don't forget to check out our handy mobile app. Just search for Troutman Pepper in your app store, download it and give it a try. And, of course, stay tuned for a great new episode of this podcast every Thursday afternoon. Thank you all for listening.