Business Better (Season 3, Episode 9): Cyber Adviser – Artificial Intelligence: An Overview of the U.S. and EU Regulatory Landscape

Speakers: Phil Yannella, Greg Szewczyk, John Kerkorian, and Timothy Dickens

Steve Burkhart:

Welcome to Business Better, a podcast designed to help businesses navigate the new norm. I'm your host, Steve Burkhart. After a long career at global consumer products company, BIC, where I served as vice president of administration, general counsel and secretary, I'm now special counsel in the litigation department at Ballard Spahr, a law firm of clients across industries and throughout the country.

This episode is part of our Cyber Advisor series where we discuss emerging issues in the world of data privacy and security. The emergence of tools like ChatGPT has demonstrated the tremendous business potential for artificial intelligence. At the same time, businesses need to be aware of the growing patchwork of laws and regulations in the US and EU, governing the development and use of AI.

Our lawyers will provide an overview of the current regulatory landscape for AI in the US and EU, and identify some best practices for businesses to employ as they consider use of AI tools. Phil Yannella and Greg Szewczyk, co-leaders of Ballard Spahr's Privacy and Data Security group, host the discussion. John Kerkorian and Tim Dickens, both of whom practice in Ballard's Commercial Litigation and Dispute Resolution group, participate in the conversation.

Phil Yannella:

Good afternoon, everyone, and welcome to Ballard Spahr's Privacy and Data Security webinar series. In our last installment, we previewed major privacy and data security developments that we're expecting in 2023. And in today's episode, we're going to focus on one of those issues, namely, increased regulation of artificial intelligence.

AI has been a source of business, consumer, and regulatory interest for some time now, but it has really accelerated over the last six months. One of the drivers for increased scrutiny of AI has been the emergence of new AI technologies that have garnered significant media scrutiny. This is well illustrated in the recent launch of ChatGPT, A powerful form of generative AI that is being integrated in chat bots with amazing and somewhat concerning results. But AI goes well beyond ChatGPT. It's already built into many layers of our daily lives from things as familiar to us as Alexa and other personal digital assistants, connected dashboards in our cars, and smart homes.

The wide range of AI technologies, in fact, is one of the reasons why we haven't seen a comprehensive legal approach to regulating AI. There is no primary regulator in the space in the United States. Different regulators instead have offered guidance on particular kinds of AI likely to cause consumer harm or lead to discriminatory outcomes. But none of the laws or regulations on the books to-date are terribly granular, at least not at this stage, which is a second hallmark of AI regulation, a lack of clear guidance.

Our goal today is to provide an overview of the most significant AI regulations on the book or in the planning stages, both in the United States and the EU. Discuss what those regulations presently require and identify particular guardrails that may exist for the development of AI.

Joining me today are three privacy and data security attorneys, beginning with Greg Szewczyk, who's my fellow co-chair of the Privacy and Data Security group. Also joining us are John Kerkorian, who's a litigator and partner in our Phoenix office, and Tim Dickens, who is based in our Philadelphia office.

Here's an overview of our presentation this afternoon. Tim is going to begin things off with a discussion of the NIST and FTC guidance on artificial intelligence. Greg is then going to talk about AI rulemaking in California in Colorado. I'm going to talk a little bit about the EU AI Act, which may be the most comprehensive proposed regulation of AI in the world. John is

going to talk about AI rulemaking in the financial services industry. Tim will handle AI in the employment space. Greg is then going to talk about potential other potential privacy laws that may be implicated by the use of AI. And I'll wrap up with a discussion of best practices. So, without further ado, let me hand things off to Tim to discuss the new NIST and FTC proposed AI guidance. Tim.

Timothy Dickens:

Thanks a lot, Phil. I'm going to start off, as Phil mentioned, by talking about how the FTC's recent AI guidance aligns with NIST's Artificial Intelligence Risk Management Framework. The FTC's recent AI guidance really comes down to two main points. First, businesses need to ensure that their implementation and marketing of AI tools is fair and transparent. Second, businesses need to conduct risk analyses to identify and mitigate any latent risks associated with any AI tools they choose to implement.

The first point, fairness and transparency, is really the bread and butter of the FTC's UDAAP enforcement power. Businesses need to be particularly wary of this issue as it pertains to AI because there is ambiguity as to the exact definition and scope of the term. Given the recent fervor around AI, including ChatGPT, like Phil mentioned, there's a significant risk that marketers try to use AI or related keywords or buzzwords to describe products in a manner that the FTC would consider misleading or not sufficiently supported by fact.

The FTC dedicated a recent post to this issue highlighting questions it's starting to ask about AI-related claims that businesses are making, but really the issues here are, have you clearly and accurately disclosed how AI is integrated into your tools? And do you understand the potential risks that could result from this AI implementation?

This leads to the second point that the FTC is really focusing on, which is that businesses must establish a framework for conducting risk assessments to identify and attempt to mitigate any latent biases associated with the AI. The importance of this issue has really become clear over the past few years in light of incidents where medical and employment-related AI tools have processed data in a manner that inadvertently discriminated against minority groups. For example, a 2020 report published by the Journal of the American Medical Informatics Association highlighted the latent risks of implementing AI as a means for efficiently allocating resources like hospital beds and ventilators during the COVID pandemic.

Because minority groups suffered a disproportionately high mortality rate during the pandemic, an AI without proper guardrails, would be at risk of allocating resources away from minority patients based primarily on their membership in these minority groups. This is obviously a bias that we want to avoid and a similar type of bias would be something the FTC would look towards as being unfair.

Now, given the nature of AI, AI risk assessments are not as straightforward as other types of risk assessments, but NIST has and is continuing to develop a framework for identifying and controlling these risks. The FTC frequently points to this and other independent frameworks as useful tools for controlling risks associated with AI.

I'll keep it at a pretty high level, but the key issue here is that with AI, a risk assessment is not a one and done situation. The business will have to conduct an initial risk assessment of the AI itself as well as any data that the AI is expected to process or learn from. In the example of the medical AI we discussed earlier, you want to look at biases that are inherent in the data and that may be associated to unfair associations.

So, additionally, because AIs learn, this review process will have to be ongoing and scheduled at periodic intervals to catch any biases as they may develop. Ultimately, the message from the FTC and independent organizations like NIST is clear. If you're going to implement AI, you need to conduct an initial assessment of the AI tool and the data it'll be processing. Identify and manage risks upfront and on an ongoing basis, and be transparent about how you're using AI tools and the results of your processing. And I'll hand it off to Greg for the next section.

Greg Szewczyk:

Thanks, Tim. So, I'm going to talk a little bit about how the use of AI in new products or on its own could have implications under state privacy laws. And before I get to what is in the regulations, I think it's important to just lay a little bit of the framework for how AI gets implicated in these laws.

What I don't have up on the screen is really the threshold question, which is, under US state privacy laws, there are certain applicability thresholds. Under California that includes a pure monetary amount, but it also includes some volume thresholds for the amount of data that you're collecting. The non-California privacy laws, which are currently Connecticut, Colorado, Utah, and Virginia, with Iowa coming on the heels, all have volume thresholds. Now, companies may not be collecting a lot of personal data, but if they start incorporating AI into some of their products, it's possible that they maybe start collecting significant volumes more without really thinking about it.

So, the use of AI in products is something that you need to consider when you're even thinking about, "Do these laws apply?" If they do apply, the use of AI in products, it could trigger some additional disclosure obligations, whether it's in your privacy policy or at the point of collection. It could also potentially constitute a sale of information or targeted advertising, which would require an opt-out right.

Now, the way this would work is if you would embed ChatGPT into your technology and you are feeding personal data into ChatGPT, since it is opensource, there's no cost, but you would be receiving something in return, which is additional data, additional technological power for your product. And these state privacy laws have fairly broad definitions of sales. So, the use of a ChatGPT embedded within your own product could end up triggering the obligation to treat that as a sale of personal data to ChatGPT or the other generative AI. If that's the case, then you're likely going to need to allow individuals the right to opt out of that process.

Another way that we see state privacy laws interacting with AI is through provisions related to profiling or automated decision-making. This is not an every state law, but it is in several of them. And the provision of opt-out rights and the performance of data protection impact assessments are triggered if it ends up constituting profiling or automated decision-making. And we'll talk about that a little more on the rules.

Under the current state law regulations, the only state that has issued regulations relating specifically to profiling automated decision-making is Colorado. And those rules were recently made final. California is going to be issuing regulations, but the set that they submitted to the OAL in February did not include anything related to these topics specifically, but they are in the pre-rulemaking phase and while they will likely differ from Colorado in some ways, there has been an effort to include interoperability between the California and Colorado rules. So, we do expect there to be at least some overlap in the approach that they take.

Now, we have up on the screen four different areas that are triggered if there is any profiling through the automated means, and that's transparency, an opt-out right, consent, and DPAs, which is that data protection assessment, rather than a data processing addenda. So, the first thing to point out is when you look at the opt-out column, there are a couple different categories here, and the Colorado approach takes different obligations as to whether or not the profiling is solely automated, human reviewed, or has human involved decision-making.

For all of them, there's still the transparency requirement, and that transparency requirement comes in the form of a privacy policy. It will have to describe what the profiling is, what decision it's going towards, what data is used, what logic is being used, why that profiling is relevant. And then there are specific additional obligations if it is going towards a product or a service that implicates housing, employment, finance, and lending.

Now, this in and of itself could just be seen as an additional burden, but when you're using an opensource AI tool such as ChatGPT, you may not really understand exactly what the logic is behind it or how the algorithm works. You may not even fully understand what data is being used beyond what you're allowing to be put into it. So, this privacy policy obligation may not be as easy as it first seems.

The next question is the opt-out. If it is solely automated or human reviewed you have to honor an opt-out. If it is human involved, there is some discretion subject to other provisions. So, whether or not you could incorporate a opensource AI tool such as ChatGPT into your product and constitute for the human involved category seems unlikely, but it is a qualification that every company should have to assess if they're going to do this. If it is not going to be human involved, and instead will be considered solely automated or human reviewed, then the ability to opt out has to be included. And so, that's something that should be considered when you're actually designing how to embed this opensource AI software into a product.

Consent comes into play after an individual has opted out. The specific consent mechanisms and disclosures are very specific for what needs to be done under the Colorado rules. And so, that is something that companies will have to consider as to how they would provide that consent.

And finally, the DPAs. Under the GDPR, this is considered a DPIA. Under California, it's going to be called a privacy risk assessment. But it is essentially profiling using automated means will trigger what is called a high-risk processing that will require an additional exercise to balance the costs and benefits of whether or not it should be done. And even beyond just the standard DPA provisions, the Colorado rules treat profiling DPAs a little different than the rest of them.

So, the takeaways are both that one, incorporating an opensource software such as ChatGPT is likely going to have significant regulatory obligations under state privacy laws, especially Colorado and California. But I think we can also see through the Colorado rules that this is an issue that's very much on the radar of the regulators. So, while it might be appealing to work this into your product as quickly as possible, it would really be prudent to consider whether or not you're going to be subject to these laws and regulations and how that might impact you downstream. And with that, I'll turn it over to Phil.

Phil Yannella:

All right. Thanks, Greg. So, I'm going to talk a little bit now about the EU Artificial Intelligence Act, which may be the most comprehensive proposal in the world to regulate artificial intelligence. It is to AI what the GDPR is to privacy, a highly prescriptive, groundbreaking law that if implemented would significantly affect companies both in the EU and outside the European Union, and likely become a template for similar regulation worldwide.

The EU AI Act was first proposed in 2021 by the EU Commission. In December 2022, the EU Council approved a version of the AI Act. The current version is now being reviewed by the EU Parliament. Any amendments that parliament makes will then have to be reviewed and approved by the EU Council. Once the law is finalized, the EU AI Act will become effective 36 months later.

Like the GDPR, the EU AI Act has extraterritorial reach, meaning that it would apply to providers that place artificial intelligence on the market in the EU, or users of AI systems that are based in the European Union or based outside the European Union but generate outputs that are being used in the European Union.

The fines under the AI Act are very high, they're even higher than the GDPR, with maximum fines of up to 30 million euros or 6% of worldwide annual turnover for violations. Here's some hallmarks of the EU AI Act. The law governs AI systems, which include both standalone systems as well as systems used as a component of a product. The current version of the act would classify AI systems into five categories, prohibited AI, high-risk AI, low-risk AI, minimal risk, and general purpose act. We'll dig down on this a little bit more.

The draft act would ban all artificial intelligence that deploys harmful or manipulative subliminal techniques, AI that exploits specific vulnerable groups, AI that is used by public authorities for social scoring, and AI that deploys realtime, remote biometric identification systems in publicly available spaces for law enforcement purposes. Again, under the proposed law, all of that would be banned in the European Union.

The law would provide for different requirements for high-risk AI, most specifically, all high-risk artificial intelligence would require a third-party assessment. The kinds of high-risk AI that the law currently identifies are AI that is involved in remote biometric identification, AI that regulates road traffic, gas, heating, AI that determines access to education or assists with accessing students. AI that is used to recruit or terminate employees. AI that determines eligibility for benefits. AI that evaluates credit worthiness. And AI that dispatches emergency services.

One thing to note is that many of the banned or high-risk systems involve uses of AI that have garnered significant concern among US regulators. Suggesting that the AI Act could become a model for future US regulation, much as the GDPR became a model for US privacy laws.

Now, a companion proposal called the AI Liability Directive is also making its way through the EU Parliament. Now, this directive would create rules to determine liability for damages caused by artificial intelligence. It attempts to address probably the two thorniest issues that are likely to arise in any civil claims involving AI. And those issues are causation and transparency. The AI Liability Directive addresses causation by allowing national courts to presume that AI caused damages if the claimant can prove that someone was harmed through a violation of a duty of care for artificial intelligence, that the output of the AI

system gave rise to the damages, and if they can prove it is reasonably likely that the failure of the care owed by the defendant caused the output of the artificial intelligence. The Liability Directive addresses transparency by allowing national courts to order discovery of information about high-risk AI to support civil claims.

Now, we're still likely many months, perhaps even years away from the EU Parliament and Council agreeing on final terms of the AI Act and the AI Liability Directive. However, this is a really important law for US companies to follow. As I've mentioned, it is extremely onerous. There are very significant fines. It has extraterritorial application and some very granular requirements. It's really critical for US companies that are either developing AI or thinking about using AI tools, to keep close tabs on this law, as it's likely to be as impactful for the development of AI as the GDPR was for privacy. So, at this point, I'm going to hand things off to John, to talk a little bit about regulation of AI in the financial services space. John?

John Kerkorian:

Yeah, thanks, Phil. Hello, everyone. Good to be with you today. I want to cover three things, how AI is being used in the financial industry today, the issues and risks that are receiving the most attention from industry observers and regulators, and where the regulatory focus is today.

First, with respect to uses, artificial intelligence is being deployed by financial institutions and showing promise in several areas, notably flagging unusual, suspicious or anomalous transactions for fraud detection and financial crime monitoring. It's also being used to personalize customer service. For example, the chat bots that are used to automate routine customer transactions. It's used in targeted marketing to bring useful products and services to the right customers. And of course, it's being used in credit decision-making.

This has the potential to allow more accurate, faster underwriting, as well as to expand credit for consumers and small businesses that may not have obtained credit under traditional underwriting approaches. But there are significant risks here that we need to talk about, and I'll get to that in a second. Just want to mention other uses of AI in financial services industry or risk management, and in fortifying internal controls, as well as, of course, detecting cyber threats and malicious activity.

As I mentioned, there are risks. In a recent request for information issued by the FDIC, the CFPD, and other federal agencies, several areas were flagged for further comment by financial institutions and other interested parties. I wanted to discuss a few of those with you today, because I think they offer the best clues about the areas where regulators will be paying attention.

The first is what's known as explainability. This refers to how an AI approach uses inputs to produce outputs. A lack of explainability can inhibit a financial institution's understanding of the soundness of an AI approach, which can increase uncertainty around the AI's reliability. Inadequate explainability can also make independent reviews and audits more difficult, and this, in turn, of course, would complicate compliance with consumer protection laws.

Now, in this area, as the banking industry noted in its comments, different stakeholders require different types of explanations based on the context. So, there can't really be a one-size-fits-all approach to explainability. I think it's a complex area. We're going to see regulatory oversight, for sure, and a balancing of interests in this area.

The second area I want to talk about and risk factor is with respect to data usage. The concern here is that inaccuracies or biases in the datasets may be perpetuated or even amplified by AI. The question regulators are focused on is, how do institutions manage risks related to data quality and data processing? And are there specific AI use cases where alternative data are effective?

Dynamic updating is the next area. This is where AI has the capacity to update and alter its approach on its own. The primary risks here concern monitoring and validating results and outputs over time to make sure the system is operating as originally intended. Oversight of third parties is another key area. Financial institutions often use third-party vendor applications and data to scale specific processes and optimize operations where they may not have the resources internally or the skillset internally to do so.

Now, I think most if not at all, financial institutions, make a conscious effort to decide whether to use a particular vendor based on their willingness to cooperate with the financial institution's risk control measures. But regulators will surely be monitoring this and will want banks to continue to conduct due diligence and perform model validation when it comes to third-party providers.

Mentioned fair lending previously. Here we're talking about the ability of AI-based credit decisions to comport with fair lending laws. No question, AI brings the potential for innovation in credit underwriting, and in fact, reducing reliance on human judgment may actually promote fair lending and reduce bias in credit decisions. However, there is no question that the agencies are focused on fair lending and the potentially discriminatory impacts of AI credit decisions. This was by far the area where regulators were asking the most questions in the recently released RFI.

The CFPB has expressly confirmed that federal consumer financial protection laws will be enforced regardless of the technology used by creditors. For example, the Equal Credit Opportunity Act and its implementing Regulation B require creditors to provide specific and accurate reasons for adverse credit actions. The CFPB has made it clear that creditors cannot justify non-compliance with these rules based on the fact that their technology that they're using for credit decisions is either too complicated or too opaque, or too new to provide that type of information.

And similarly, as Tim mentioned, the FTC has stated that it is evaluating fair lending practices under its broad Section 5 powers to ensure racially biased algorithms are not deployed, and it's issued a warning that it will hold businesses accountable for credit discrimination. So, in the fair lending area, it's critically important for financial institutions to engage in free implementation, fair lending testing, ongoing monitoring, and periodic back testing of model outcomes and trend analyses.

The reality is that, Phil mentioned this, for the time being, financial institutions are left with a patchwork of guidance, and it makes it challenging for financial institutions to fit new technologies into that older guidance. My sense is institutions are being appropriately cautious about the deployment of AI and probably will continue to be until there is more clear or comprehensive regulatory guidance. As mentioned, there are states, California, Colorado, that are more active here. We'll be watching that. We'll be watching the EU and other foreign jurisdictions to see how that shapes the regulatory environment for financial institutions. So, with that, I will pass the baton.

Timothy Dickens:

Thanks a lot. The use of AI in employment implicates many of the same issues we've been discussing throughout this presentation. One of the thornier ones is one which Greg addressed more broadly, and that's compliance with the patchwork of state privacy and employment laws.

So, as Greg mentioned earlier, as of January 1st, the California Consumer Privacy Act now applies to employees. This means that any qualifying businesses that use AI-related technologies to hire, monitor, or otherwise process the personal information of California employees, must provide appropriate disclosures, enter into appropriate data processing agreements, and potentially provide users with opt-out rights to the extent that the tool may constitute a sale of information under California law. As the California regulations develop, businesses will likely also be required to conduct a data protection impact assessment prior to implementing any employment-related AI in California.

Going beyond the CCPA, use of employee monitoring tools will likely trigger state employment laws such as Connecticut, Delaware, and New York's electronic employee monitoring laws. While these laws are not AI-specific, AI tools used to analyze or review employee conduct or actions would likely trigger their notice and employee acknowledgement requirements. This may go beyond minimum privacy policy disclosures and require that employees specifically acknowledge that they are being monitored and that it may be used for given purposes. That's going to depend somewhat on the state law at issue.

So, additionally, there has been an uptick in state and local laws designed to control the use of AI tools specifically in the hiring process. Again, these focus pretty heavily on transparency and controlling biases. For example, the New York City Council recently enacted a law requiring that businesses that implement AI or automated decision-making tools for hiring, conduct a bias analysis and expressly disclose the use of such tools.

Any large companies with employees across multiple jurisdictions will have to stay on top of this patchwork because it's really actively developing, and it can vary. In the case of New York, it's at the city level. This could be a state issue, a city issue. It's something that's got a pretty broad base that companies are going to have to stay on top of.

In addition to the state patchwork, federal agencies such as the Equal Employment Opportunity Commission have published initiatives and guidance regarding the use of AI in the employment context. This guidance is largely aimed at fostering the implementation of AI while ensuring that it is not used in a manner that results in unlawful discrimination. For example, the

EEOC and the Department of Justice have both published technical guidelines for proper implementation of AI tools. The EEOC focusing on ensuring compliance with the ADA, and the DOJ focusing on the public sector.

Importantly, the EEOC has already brought a suit alleging that a business violated federal law by implementing an AI recruiting technology that discriminated against older applicants based on their age and gender. We expect to see more of these discrimination type issues arising as AI continues to proliferate.

Finally, as Phil mentioned, businesses operating in Europe or the UK should be cautious of implementing AI tools in a manner that could result in hiring, firing, or compensation determinations. Such determinations may constitute automated decision-making, raising additional disclosure, consent, and opt-out requirements under the GDPR. And that's even before we get to the AI initiative that Phil discussed. Ultimately, businesses looking to incorporate AI into the employment and recruitment context should use extra caution to ensure compliance with this complicated and evolving legal landscape. And with that, I'll hand it off to Phil and Greg.

Greg Szewczyk:

Thanks, Tim. I'm going to move through some of these relatively quickly, but I think one thing we want to make sure to highlight is that on top of some of these bigger areas that we've been discussing so far today, there are also several other legal areas that could provide potential landmines for any company that's looking to incorporate opensource AI or any other kind of generative AI into their products.

The first one that I want to talk about is state wiretap laws. For those who have been paying attention over the last year or so, we've seen a flood of wiretapping class action lawsuits in the US, that are based on state wiretap laws that are two-party consent states such as Pennsylvania and California. The model of these class action lawsuits that we have seen so far typically focuses on the use of website analytics software such as session replay tools, tracking cookies, and chat bots. We've also seen it in the healthcare context with patient portals, and we've seen it with insurance quote tools.

The general theme of these lawsuits is that these third-party tools capture communications between the first-party website user in realtime, which plaintiffs allege is an interception of these state wiretap laws. For the most part, what we've seen with ChatGPT so far is an interaction with the bot itself. So, it would be a first-party communication. But if you start embedding ChatGPT or other generative AI into a product, there could be an argument, similar to what we have seen in the website analytic context, which is essentially that there is a third party who is intercepting ... And the third party would be the opensource AI, that's intercepting a communication without the consent of the user.

Now, the reason these lawsuits are so popular is that they provide a private right of action with statutory damages. So, businesses that would be incorporated in this should be cognizant of these potential types of lawsuits that can include one to $5,000 per violation. And when you think of the number of website users, or your product users, that can add up quickly.

Another area where we have seen certain types of potential liability that could be applicable in the opensource AI context is web scraping, which would likely come in the breach of contract context. What we're talking about here is if you are using generative AI or opensource AI in a manner that lifts information from other websites, depending on the websites that is going to be used, and there may be no way in which you could realistically prohibit it from being used or collected in that manner, it could violate the underlying website's terms of use. And we've seen this come up a lot in the data broker context, and there has been a distinction in the case based on whether or not the tool being used to scrape the data is an automated tool.

So, the companies that have been using this so far have been fairly keyed into this kind of case law and have been doing their best to try to stay out of the crosshairs. But for companies who are just going to be incorporating what appears to be a free tool into their product, it might not be front of mind.

Another area where we expect to see some potential regulation, most likely, would be in the healthcare context through the 21st Century Cures Act. The 21st Century Cures Act was signed into law in 2016, and it provides the FDA with authority over digital health products, including on the cybersecurity front and over AI. We have seen the FDA be active on this front already as there have been a variety of studies that have shown that the use of AI in healthcare products can lead to discriminatory effects. But the use of AI in healthcare products can also be an incredibly powerful tool, and so, is becoming more and more

pervasive. To the extent that your company is looking to incorporate AI into either the provision of healthcare or a health tech product, need to be particularly cognizant of potential risks under the Cures Act and under the FDA's guidance.

Another area is products liability. Under the caseload that has developed over the last several years, it is likely that AI itself won't constitute a product, because it's not tangible personal property, distributed commercially for use or consumption. Instead, it will likely be considered a service, but that still needs to be decided. But even if it is considered a service and not a product, the product that AI would be incorporated into may still be a product subject to strict liability. So, knowing how that AI is going to interact with and control the product is going to be very important, especially in light of the limitations on liability and indemnification that are baked into the terms for opensource AI.

And finally, intellectual property. Now, the concerns on this front can be broad, but a very clear one is going back to the web scraping context, which is, if you are using opensource AI or generative AI to collect broad amounts of information from the web, and then it's going to be reproduced in any way, you need to be cognizant of potential trademark, copyright, and other intellectual property violations that could be out there, and what safeguards you may be able to put in place to prevent any liability on that front. So, with all of this information, I'm going to turn over Phil to try to give some high-level pointers on what companies should be doing right now.

Phil Yannella:

All right. Thanks, Greg. So, we're going to wrap up our discussion today by talking a little bit about best practices for the development or use of AI. What can companies do if they're either contemplating developing or using AI, to protect themselves from liability? Well, first and foremost, companies need to track AI regulation. We've gone over a number of different laws proposed in actual, we've outlined some risks for specific industries, but the regulatory landscape is not likely to stay static in this area. We're almost certainly going to see new laws proposed here in the US and in the EU to govern AI. It may be comprehensive privacy laws like Colorado or California, that also cover AI, but there also are likely to be AI-specific laws. So, companies need to track those carefully. We offer a number of resources for clients who want to track proposed regulation, including our blog and webinar series, as well as legislative trackers that we can provide to clients.

A second step for businesses considering use of AI is to develop policies that ensure managerial ownership of AI development, implementation. Historically, regulatory oversight of emerging technologies has focused, at least initially, on the absence of managerial oversight. We saw this, for example, in a case of data security. One way businesses can manage legal risk arising from AI is by incorporating AI technology into a company's written information security or privacy program.

Another way to achieve this goal is to develop a separate AI governance plan. This type of policy would provide guiding principles and actionable requirements for the development, purchase, and use of AI. Businesses looking to use AI products should also begin ensuring that the issues are raised at the board and senior management level, whether it's part of an annual privacy assessment or otherwise. There is a clear regulatory consensus across industries that boards and senior management need to be actively involved in data privacy and security, and I think we can expect to see something very similar for AI.

A third step for companies to mitigate their risk is to begin to focus on a range of vendor management challenges that are likely to occur when you're implementing AI. Vendor contracting is going to be challenging, particularly for businesses that are subject to both US and EU privacy laws, which require data processing agreements that in turn, in part on the characterization of the vendor is either a service provider or a business/controller. And whereas traditional vendors are typically classified as processors or service providers, AI vendors may not, because many of them will be holding onto the data for their own purposes, meaning they'll be holding onto data to train the algorithms.

Businesses should also be careful to understand and control the scope of automated processing by opensource vendors. For example, a domestic company that uses opensource AI tools to scrape the internet for specific categories of personal data, might find themselves triggering the extraterritorial application of the GDPR, which would create new privacy risks and obligations for a business.

Conversely, European businesses should consider the potential risk of transferring data to AI vendors in the US, because that may be subject to government access requests through FISA warrants. So, depending on the nature of the data being transferred, European businesses may need to implement additional safeguards to lawfully transfer personal data to US-based

AI vendors. To identify and mitigate those risks, businesses are really going to have to understand and contractually control the nature of the processing at issue.

So, that brings us to the close of our webcast today. Before we leave, one last word, in addition to speaking and blogging about cyber litigation, we've written a book on it, and it's called Cyber Litigation by Thomson Reuters. It covers a wide range of data breach, data privacy, and digital rights litigation, everything from retail data breach litigation to online tracking litigation, to website accessibility claims. You can check it out at https://store.legal.thomsonreuters.com/law-products/Treatises/Cyber-Litigation-Data-Breach-Data-Privacy--Digital-Rights-2021-ed/p/106731568. Again, it's published by Thomson Reuters. Thanks, everyone, for joining us today, and we'll talk again next month.

Steve Burkhart:

Thanks again to Phil Yannella, Greg Szewczyk, John Kerkorian, and Tim Dickens. Make sure to visit our website, www.ballardspahr.com, where you can find the latest news and guidance from our attorneys.

Subscribe to the show in Apple Podcasts, Google Play, Spotify, or your favorite podcast platform. If you have any questions or suggestions for the show, please email podcast@ballardspahr.com. Stay tuned for a new episode coming soon. Thank you for listening.