Welcome. You're ready for today's webcast and so are we. Please note that this webcast is being recorded and the on demand replay will be made available via e-mail along with a download of the slides in a few business days following the webcast. For the best webcast experience, we recommend viewing today's event via laptop or desktop while using Google Chrome or Microsoft Edge. Please ensure you have a reliable Internet connection, all other applications are closed and that your pop up blocker software has been disabled. We want today's webcast to be as interactive as possible, so please send us your thoughts and questions throughout the presentation.

During the live broadcast, we will answer as many questions as time allows. Next to the slide window, you can submit questions by typing them in the Q&A box and hit submit in the Resource List widget on the Dock. You can access today's content related handouts. To stay updated on all of our future programs. Sign up for our webcast digest today by visiting the Call to Action window on your screen or by visiting the Resource List widget on the Dock. If you experience any technical issues viewing or hearing this presentation, slides not advancing, or the inability to answer polling questions, please refresh your browser window.

If you continue to experience issues, click on the question mark help widget on the dock at the bottom of your presentation window. Please note, if you're viewing today's webcast in full screen mode, you'll not be able to see the pop up polling questions. If a polling question is announced and you cannot see it, please be sure to exit full screen mode. And now please welcome our host for today's webcast. All right, hello, and welcome to our webinar today on identifying and securing against the risks of generative AI. My name is Josue Berra, I'm a manager in GT Cyber and Privacy practice, and I'm looking forward to spending some time with you today to discuss generative AI.

We'll talk about what it is, how it's used, and what its risks can be. Then we'll get into some approaches about how you can secure against these risks utilizing Microsoft security tools you might already have access to. To support me in this discussion, I'm lucky to have two cyber experts joining me to provide their perspective and experience working with organizations on this very topic. First off, Mike Lord from Microsoft. How's it going, Mike? Hey, thank you. It's going great. Yes, my name is Michael Lord.

I'm a data security and compliance global black belt for Microsoft. I've been here about 12 years. I've been at Microsoft and prior to that time I worked for a large telecommunications company building cybersecurity appliances. Happy to be. Here. Well, thank. Well, thank you for being here. We're also joined by my boss here at GT, John Pierce. John, can you give us a little intro? Sure.

Great to have peace speaking with everyone today. My name is John Pierce. I'm the principal of Grant Thornton in our Cybersecurity and Privacy group based in the Washington DC area. Been at Grant Thornton for almost 8 years now and in this industry for for close to 25 and look forward to have a good discussion about this topic today. One thing I'd like to remind the audience of is even though this is a non CPE webcast, you do have the ability to ask questions in real time. Please go ahead and submit those as you can and we will answer those as we can through the presentation. And if we don't get to them during the presentation, we definitely will at the end.

Bert All right, so before we get started with everything, I'm just going to set a quick agenda. First off, we're going to start off with the basics. What is Jenna AI and why are we bothering talking or thinking about it? We'll discuss its risks and what you may not be considering when deploying AI systems. Next, we'll address some of those concerns that we brought up and how we can potentially use Microsoft technologies, specifically Microsoft Purview and its AI Hub feature to secure against these risks. Then we'll get into scenarios you might find yourself in after deploying an AI system.

We'll run through questions you may get from leadership or from stakeholders and approaches you might take to, you know, away. Lay these concerns and talk about the risks. And then as John mentioned, we'll have some time for questions at the end. But again, if you have anything come up during the presentation that don't

hesitate to include it in the Q&A, we'll try and address them as we go along as well. So let's get into the material and let's start off with the the big question of what is generative AI? Well, generative AI is a type of artificial intelligence technology that can respond to prompts to produce various types of content such as text, images, audio, and video.

Even likely already had some experience with Gen. AI. If you've gone on social media, you can't go a minute without seeing a crazy or maybe more disturbingly a little less crazy image and video simulated or created by Jen AI. You may have heard your kids are maybe using it to help out with their homework, or you may have heard that one of your Co workers using it to, you know, while working from home or at the office. And you may be familiar with some names of major players such as Jet GPT, Microsoft Copilot, new journey, Gemini. That's a lot of different varieties out there.

So there's a lot that you can do with generative AI and it's not just for laughs and social media business applications for Gen. AI usage are growing. And again, you may have already had some experience with them, whether you've recognized it or not on the screen. We have a few examples. It's certainly not limited to what we have here though. But personally, I've used JNII in many of these use cases that we have to scratch. My job's a lot of reading reports, referring to notes taken during meetings, and basically slogging through data and details. AI note capture is a real blessing and asking AI to look through these notes and put the bottom line up front is a great, great time and resource for for saving you time when going through the literature.

Using generative AI to help generate code is a great resource as well for getting started for a project or thinking through a problem. And especially if maybe your, your programmer hat like is it's a little dusty like mine doing a lot of these assessments, but it's a great resource to again, talk through a problem when analyzing data sets and, and looking at data Gen. AI can be a really great resource because you can talk through your problem using natural language planks like Copilot. You can describe a problem. I'm looking for something within a data set rather than having to figure out the query that you're looking for.

This natural language capability makes it also great for communicating. And so it makes it a great chat bot or training tool to help guide users through training or, or learning any sort of use case. And so as I mentioned, this aren't the only use cases. So, Mike and John, what are some general use cases that you may be encountered while working with clients? Appreciate it as well. I'll give a couple of examples and, you know, pass the microphone over to Mike. I think where people are probably, you know, most excited is probably anywhere they could see, you know, business process automation, right, right.

Really simplifying automating, you know, simple tasks, routine tasks and workflows. I think that's why you see a lot of, you know, business units get very excited about and there's, you know, a number of different, you know, kinds of, you know, use cases, you know, some in the fraud detection space where we see, you know, financial institutions, you know, leveraging Gen. AI to find those unusual patterns and, and potential fraud as well. So I think the capabilities are are really limitless and really despite depends on on the area of work Mike and you'd like to add to that.

Yeah, I think one of the the major use cases that we're seeing, at least from a Microsoft perspective, is a lot of users are really excited to use generative AI in conjunction with meetings to help them manage meetings. That can be everything from generating transcripts to using AI to do translations and summarizations, as well as do things like capture action items and effectively make a lot of those repeated tasks, A lot of those things like meetings and emails and document summarization, those repetitive tasks that you're referring to, John, and really making the ability to digest that information so much faster. So when Microsoft talks about AI, we tend to call it copilot, not autopilot.

It's not there necessarily to do something for you, but they're to help assist you do something. So one of the jokes that I always tell about a generative AI is really the generative AI is almost like having a recent college graduate that has no expertise in your subject area come to work for you. So at that point, it's almost like an

intern. So it can do a lot of really smart things, but there's some some context and some review that you still have to do that'll help you, you know, do things faster. And given that kind of context, we're starting it to see it move into the cyber security realm into you pointed out business process automation, but any of those mission specific or business function specific capabilities where you can analyze anything and then provide a summarization back very quickly.

Thanks, John. Awesome. So I think we've all seen the potential and at GT we definitely see the potential, but it's obviously not just us. Gartner released a statement saying that more than 80% of enterprises will have used generative AIAPIS or developed generative AI enabled applications by 2026. This means that it's going to be something that you're going to need to think about. And if you're like me, you're probably already thinking about the risks. And if you're not like me, you're in this webinar, so I'm going to tell you about the risks anyways.

The first risk comes with how AI is used. We've talked about some of what you can do with Gen. AI, but again, as we mentioned, we can't possibly cover it all. And there's a lot of different products out there and they aren't all the same off with the same security features. And that can happen even if they come from the same vendor. The Next up is where is the data going? Generative AI requires an input to produce an output. That picture of George Washington riding a bald eagle required a prompt.

That code analysis required some code to analyze. Is it all right that your employees are uploading data to these products and services? Do you know what that data is and are you aware of any rules that you might be running into issues with? So let's get a little bit deeper into these risks to explore that topic. So how is AI being used? As mentioned, you may have already had some experience with Gen. AI out of the business context. Jet GPT Co, Violet and other tools are really popular with the general public, and these tools offer a lot of flexibility.

And there's a lot of use cases that the developers of these tools and the business leaders who are marketing them probably never considered. But they're out to the public and they're in the public hands. So the public going to be creating them. And it's hard to protect against something that you didn't even consider somebody might do. It also means that these tools are going to be coming a reading, transferring from novelties to basically expected tools, and you'll probably going to need to figure out how to accommodate them. If a new generation is used to these tools, using these tools during their education, in their training, they're going to be harder to say no to when they're entering the workforce.

And will you be able to take advantage of the efficiencies that these products can bring? And can you afford not to? At the same time, you don't want your data, especially sensitive data, to be exposed through an insecure service. And to complicate matters, your employees might be demanding the usage of one of these tools. So you go out and buy one of the tools and you buy the enterprise version, which comes with more controls over how that sensitive data is processed and other security controls that you can implement. So then you find out that your employees are signing into the tool with their personal e-mail and accessing the free version.

So Mike and John, is there any more where we can tell the audience about what you're hearing from concerns and from business and IT leaders about the usage of AI tools? Sure. Yeah, maybe, maybe I'll start, Mike. I think the biggest one is I think many business leaders are worried and concerned that they're not using AI solutions. They're going to fall behind their competitors, right? That's, you know, so realize it is a huge business driver for organizations to look at these technologies. And when I look at this, I think it's very similar to 8-9 years ago with SAS solutions as they started to come to be where we can, you know, very easily, you know, bypass organizational procurement models, use a credit card and have a capability.

So on one hand, you have, you know, businesses that are, are convinced they're going to be left behind if they don't rapidly adopt AI solutions, right? And then the other concern is, well, how do we control that as well? So it's something we have to get a hold of and get ahead of prior to that business going ahead and and you know,

taking the, you know, the the reins very quickly and leveraging those solutions. And you know, John, as as an individual who early in my Microsoft career dealt with the migration to the cloud and the creation of something called Office 365, we had a lot of conversations with IT and, and talked about IT objectives and how to take these tools out of their own data centers, their own servers and put them into the cloud.

Well, now with generative AI, what I'm starting to see is the audience has really changed. So now folks from legal other individuals who are never involved in discussion before have really come in to assess risk because of what they've read in the news, what they've seen and and some of the the stories that you see out there around generative AI and its use. So there's a lot more due diligence that I'm seeing where security teams, Sisos, as well as legal, as well as other folks around compliance are starting to ask what's happening with generative AI. So your audience that's really making this determination around the tools being used has expanded beyond your traditional IT boundaries and started pulling in all these other roles so that they can really see what's happened.

I think from that perspective, when we talk about sensitive data that can be exposed, that's really what keeps a lot of people up at night is am I complying with rules? Am I taking privacy data, and am I passing it along to the AI tool and training that AI tool when it's used? Am I taking my sensitive information, you know, my intellectual property, my, my things that make my business run sensitive contracts, am I sending that off into another region or another country? So questions like GDPR, California Privacy Act, as well as some of the other broader EUAI Act and other things really do come into play and really become part of the discussion.

So really it really has changed the type of risk discussion we have to include these other organizations that typically weren't involved in IT discussions beforehand. And it's good that we're having these discussions because that is one of our other risks getting thinking about generative AI and where the data is going. This is going to be typically a problem that's going to be more relevant for Gen. AI products. These tools typically run services that require a lot of processing power on the back end, which supports a cloud model. And so that means that when your user base is entering texts, entering files, sending over messages to these chat clients, it's data that's no longer in your environment and under your control.

Furthermore, as Mike mentioned, this data is often used to improve performance, meaning that the data that's submitted is going to be taken, is going to be looked at potentially by the developers used by that model to enhance its capabilities and its ability to respond, potentially exposing your data. So I know we touched on a little bit, but is there any other insights that you have on get exposure risk with Gen. AI Mike and John? Well, I yeah, let me chime in for that one. You know, speaking from from a Microsoft perspective, we actually have multiple products out there, multiple generative AI products. And if you just look at the Copilot name alone, we have used this as an application along a lot of different different realms.

And some of those products are targeted at consumers and our consumer products. And some of those products are in fact Enterprise Products with enterprise terms and conditions that really focus on all the compliance pieces that come with our enterprise cloud services. Things like GDPR really come to mind, but all the privacy rules and everything. And I, I really do think understanding that and understanding what those tools are and then for organizations to be able to assess that risk is the foundation behind this. Knowing that they're private data, whether that be data that they're legally required to keep private, like customer information right to be forgotten or their, their sensitive information once again, is really, they understand what the implications are of that and they understand how that AI service works.

And really that's what I spend my days doing is explaining how the different products work within their context, because organizations really want to understand that in order to adopt it and use it. And to the point that you made earlier about we're just finding that people are using these tools. Organizations are finding that no matter what they do, people are using generative AI. So now how do they funnel them into tools that really align with the business needs, the organizational needs, including things like privacy? Now, John, what kind of conversations are you having these days along those lines?

Yeah, the main question, you know, people have is where do they start and how do they protect against these risks? And I wish we had built in a, a poll question to ask the attendees how many organizations actually have a, you know, an AI or a Gen. AI usage policy at their organizations, right. So I think that's one of the main struggles that the people are having today. How do we even, you know, govern its usage? Because to your point, Mike? If you don't govern it and monitor it, I guarantee it's probably being used today in your environment, right?

And I think there's some, you know, already some very well known cases with like technology companies who didn't have, you know, a usage policy, you know, developers, you know, leveraging, you know, different solutions to, you know, check their source code and then exposing confidential information, right? So those situations can happen at any point. You know, Mike, I think your, your point around how, you know, different, you know, AI solutions are meant for different audiences, consumer versus enterprise goes back to, you know, developing, you know that Jenny, I policy how it's going to be used in the organization, you know, how it can be treated because you don't get ahead of it. You got to figure out a way to convert them to, to authorized solutions, you know, much later, right.

So I think it's a big concern for a lot of folks. It's kind of come up fairly quickly. Organizations available to adopt these solutions potentially quickly and then the risk of data exposure if using the wrong solution could be high. And I think, you know, organizations are sometimes taking a maybe let's not look approach, which there is a perspective on that. But I think that the risk of data exposure in other areas around, you know, breaching privacy regulations can be very high. And let's talk about those regulations.

They, we do have law and industry standards to consider and AI specific regulation such as laws related to employment decisions are already in place and broader regulations are being passed and being considered around the globe. And additionally, while it's maybe great in theory to try and create a, you know, huge mash up that that executive producers in Hollywood would only dream about their intellectual property, lawyers may disagree or see differently. And Gen. AI tools can generate content that closely resembles existing copyrighted work, leading to potential legal disputes or IP theft concerns. And so again, Mike and John, are there other things you've been hearing about compliance or issues with generative AI there?

Yeah. I mean, I'll, I'll start. I think you're going back to to Mike's point, you know, the biggest ones you naturally are, you know, the data privacy regulations and understanding how you know, proper consent and data usage is, is really being, you know, leveraged in those models. And I think also just with regulators around, you know, transparency and explain ability as well. I think we had a really interesting question too. And maybe Mike, we kind of weave this into our answer is, hey, is this just like how we had SAS solutions, you know, back in the day?

What is there something different here? And I think when it comes into, you know, the various language models and data usage, you probably could have more exposure there. Do you have any perspective on that as well from that question? Yeah, I think there's a couple things going on here. So you know, there was a question that came in the chat around generative value tools and SAS and drawing more light to to what has always been true for SAS. Well, that is some of those tools are correct. So I'm going to speak authoritatively from a copilot for for Microsoft 365 perspective.

That's exactly what it's doing. It's using the existing permissions of a user in SAS and it is of course then reaching out to ground that data on an individual SAS service files, what they have access to. What we're seeing is two different things along those realms. 1 is that data moves through an application. Is that application compliant with the regulations that are in play? So what the question that everyone should ask, and I can say this for Copa authoritatively for Copilot, for M365, is it maintaining my existing SAS obligations and my existing, you know, existing privacy boundaries?

Absolutely. That is not always the case with certain tools and that extends further into compliance. So what I'll see for example, is that the location where data is processed, these everybody seen the stock prices of the, the hardware manufacturers for generative AI, they are going through the roof. Well, they can only build these systems in certain places. They can only host them in certain areas. So you may be crossing geographic boundaries as you process this data. Understand a lot of vendors including Microsoft, that's easily controllable. However, you need to actively ensure that those controls are in place.

So you may for example, take information, that information might be in the United States and you're processing, processing that data in Europe and then returning answers back to the United States. Is that OK? Might be OK. It really depends on an individual, how they're aligning their business practices, what that data is, etcetera, and really driving that forward. And you know, further what we're starting to see in terms of compliance. There was a question in the chat about SoC 2 reports readily available for generative AI SAS services.

That varies. So typically your CASB vendors including Microsoft, we have Microsoft Defender for cloud apps. We're tracking that information. CASB vendors are showing those authorized and unauthorized AI apps. However, that doesn't mean that your users are using them and they're not going to something else. There's a difference between enterprise ChatGPT and ChatGPT. There's a difference between Copilot with enterprise data protection and Copilot, the consumers that use their Outlook.com Xbox account and log into and interact with that data. And you're going to want to make sure that users are using those compliant AI tools that have those reports as opposed to using the freemium service that is gathering that data and selling that data.

And it's always been a joke that, you know, if I'm not paying for this tool, then I am what's being sold, right? So there's those jokes about free e-mail services, Research Services, that they're compiling that data and selling that data, selling that marketing data. Well, the same thing is true in generative AI. So understanding whether it's a enterprise tool in an enterprise SAS service or whether it is a free service that may not have the same terms and conditions as the SAS service it's integrated with. Awesome. Well, thank you for all that.

And I think that's enough for a doom and gloom for right now, though. And while these risks may see, you know, a little daunting, there is hope. And let's hear it from Mike and Microsoft for helping us with this. Microsoft Purview is a platform that provides tools that enable information protection, data governance, risk management and compliance. And Microsoft Purview has introduced a feature called AI Hub, which unifies this functionality to target AI issues directly. So to answer our questions about how AI is being used, AI Hub provides visibility and breaks down which AI tools are being used, who's using them, and how they're being used, along with that key question, what version is being used.

With that visibility, you can determine what where your date is going, if there is issues, you can implement rules to limit access or functionality of AI tools to help limit the state exfiltration. And then finally, AI Hub provides tools to monitor emerging regulations and standards and assess how you stand in compliance. Now, Mike, is there anything else that you'd like to tell us about a development of of purview or? Yeah. So AI Hub was really for us was generated as a tool to help facilitate use of generative AI And, and a lot of discussions came up about discovering data risks in use of AI.

And there were some very early stories of in the legal realm, in the manufacturing realm of things that happened with information that was used in generative AI and individuals not necessarily understanding those tools. So from our perspective, we wanted to create a tool that allows you to assess that risk and really drives back to your ability to ultimately prioritize critical data risks and then allow data protection measures to be put in place. And I think the real key behind that is that needs to line up with your policies, right? So you have the ability to, you know, when we talk about things like European Union, workers rights, privacy, etcetera, that it allows you to basically get out front, talk to your workers councils, talk to your employees about how these tools are used and really ties into an acceptable use policy and then provided technology then can snap into this framework that you put in place.

Awesome. So now that we know something about Gen. AI and it's risks, let's go into two scenarios you may find yourself in in trying to address after deploying AAI tools in your environment. How can you monitor usage and how can you prevent data loss? Let's start with monitoring usage. So in this scenario, management has some Gen. AI usage concerns and wants to understand how and where it's being leveraged and if making further investments make sense. For a little bit of background here at your organization, generative AI has swept in and everybody across it is trying to use it and see how it will fit in with their work flows.

You've recently enabled Copilot for your organization and it's seen broad adoption. However, you are aware that your users are also potentially leveraging other unapproved AI chat clients such as ChatGPT. So this is a pretty basic scenario and plays right into the strengths of monitoring and Purview. We're not going to go into the technical details of setting up Purview during this webinar, but if you work with Purview in the past or potentially deployed Microsoft Endpoint DLP, you may find that some of the prerequisites for getting this set up are already taken care of. But going into the scenario, let's say we're in our meeting with management.

And so here's some questions that management may be asking you and here's how you can potentially respond when you're using Microsoft tools. So they may just come up and say, all right, so our problem is how can we identify what AI tools people are using in our organization? And you can say, well, we can enable Purview and we can use its AI Hub through that. We can centrally monitor AI applications, including Copilot, which we have deployed along with our third party or any third party AI tools that are being used. And so they may say, well, that's great. But.

We really are more interested in understanding how what you know, what state, what type of data is being shared with these tools and if you know, does it pose any risk to us. And so you can say that AI Hub allows us to monitor the categories and classifications of files and texts that is entered or pasted into AI tools. So you're capturing what's shared and getting an understanding of either what that content is or at a high level, what it might be classified as and classified as and identify those risks me and what might also hear. How can we monitor if our AI usage is in line with laws and regulations that are applicable to us?

And you can say that well, with AI Hub, we can have run assessments of our environment, including other cloud environments, not just limited to Azure that we have in our implemented in our organization and it will provide us recommendations on actions that we can take to comply with AI specific regulation. So that's a starting point and that addresses, you know, the main questions you may hear from one of these scenarios. Mike and John, is there anything else that you might hear from leadership questions you might expect management to ask of the usage of of AI tools and monitoring or other perspectives that you could offer?

Yeah, I think visibility is really key, right. I think something that Mike said at the beginning, you know, some of the primary use cases, you know, for, for these solutions, you know, writing emails, summarizing data, this is something that every single employee has to do and probably loads a little bit, right? So naturally the ability to have a solution to help you, you know, do so is very compelling, right? So I go back to, you know, always grounding, you know, usage within policy. I think we have some questions around example, you know, Jenny, I have policies. We do have the ability to to share some of those.

And then having the ability to actually see, you know, where these solutions are being used. You know, we've had clients who have, you know, leverage this capability in their environment and like the old saying, you turn on the light, you see the cockroaches, right? And you're going to see very quickly where people are probably leveraging the solutions and exposing data. Potentially they can cause harm, you know, from whether it's a regulatory, you know, law. So it could be, you know, exposure of confidential data and the whole like, so I think, you know, going through this kind of exercise you get a lot of visibility around usage and really challenge organization how they get ahead of it as well.

And you know, when, when you really, when you really dive into this, there is so much value in generative AI to the point where I don't know how I lived without it before this. The ability to summarize emails and the

ability to handle when I'm double booked and triple booked in a meeting really does matter. That really does make me more effective at my job. And, and you're spot on, John, with that, things you loathe to do, right, like trying to juggle so many meetings. So really just ensuring that I we do that in a safe manner is really what I what I have focused on professionally and in driving that down as a key value.

Awesome. So let's move on to preventing data loss and data protection. So as we dive into this, let me let me just bring this up. We've talked about how to find, right? Like I'm looking at it, I've got that visibility, I know what's going on. I can take that information and I can reactively look at it and determine what's happening. Now we're talking about shifting into a proactive control type of scenario. We're talking about how do I now do something about it?

This is of course optional, but really for me, the key point when we see this as we dive into this detail is you're going to use the term data loss prevention. And that really does drive at the heart of this. So please take it away. Thanks. And and yeah, so this is a great continuation of of maybe the scenario that we just talked through and we've identified what, you know, how AI is being used and management's been talked through it. And so now they want to create a policy because they've seen what people are doing that will disable employees in the farm finance department's ability to copy and paste content into generic AI websites.

And I think this is a great scenario because it shows some of the flexibility that we have with with Microsoft purview. And I've heard in particular a lot of about a lot of interest in this ability to stop pasting or uploading into these sorts of products because it's not a hard stop, but it makes you think. And if you make somebody think about the action that they're taking, it can help. Well, maybe make them make a smarter decision in terms of what they're what they're sharing. So in the background, your organization has implemented Copilot to great success and user adoption.

However, again, after reviewing the activity alerts generated in Burview, management has decided to limit this specific group of users and they determined that finance should not have the ability to copy and paste sensitive financial information into these websites. So we're in the meeting and we're scoping this out. And so management will, it's going to ask you, well, what can we do? Can we actually restrict users from copying sensitive data into AI applications? And you can say, yes, with Purview AI Hub, we can block or warn users when they attempt to paste or upload or enter sensitive information into generative AI applications using the DLP capability.

Now you may say we don't want to block, you know, that's great. But we don't want to block all usage of AI. As you know, as been mentioned a lot, it's been, it's a great resource for accelerating people's capabilities, answering emails, doing that sort of thing. But we need to control data as dictated by our policy. So what can we do? And you can say that Purview allows us to leverage sensitive sensitivity labels that we've already. Created. And encryption to ensure that the sensitive data remains protected throughout its life cycle.

This includes when that data is being used by applications like Copilot or other regenerative AI applications. So, you know, we have some controls, we have some flexibility, but management smart and they're recognizing, you know, these Gen. AI apps and the companies they're evolving. As Mike mentioned, there's different business models around it. If it's free, it's free, but it's maybe not so free. But business models change, companies change, perspective change. And how can we as an organization track these changing risks over time? And what you can say is that again, with Purview, we can automatically run periodic risk assessments of those AI applications, block applications that are risky.

And it can, with this, we can reduce the need for manual intervention while keeping our environment more secure. So again, Mike and John, is there anything else we want to touch on our questions we might hear from management around data loss in these scenarios? Yeah. I think when you look at these, you know, scenarios, you can kind of come up with different use cases and workarounds. And I go back to how Jen AI is, is very

attractive to probably every single employee in your organization, right. So can we design models by bad word to use for this design?

You know, structures to, you know, stop, I'll say by 85% of the the the good people from from doing things and maybe it can fall right. So I go back to it's always challenging to then stop a business process from happening that's already leveraging me. So the sooner you get ahead of it, you know, with. How those policies, you know, the better and I think, you know, applying, you know, some, some broad, you know, controls in place like you see here as a starting point and allows you then to to have organizations or, or people who want to use these technologies go to the proper channels.

Well, I do the risk assessments around them and approve them to be used. You know, there as well. There's always be a work around around these, right? If you, if you look at this, you say, oh, well, what about this person? Could they do this? Can they go off network? Could you do these different things? They can, I think the main thing is let's, let's get the vast majority of people who probably didn't know any better, right? They, they started doing it at home.

Maybe they, they use a solution to help them figure out their, their children's Algebra 2 homework. Maybe a certain speaker may have done that last night and then they'd come do it at work the next day, right? So we might be able to stop those people from doing those things because probably could be against policy because to Mike's point, you know, you know, you know, free means you're selling your your own data, right? So we want to stop at an organization too so crowded in policy. Let's focus on the 85% solution. You know, most people are probably leveraging these tools, not realizing there's a policy in place.

And these these solutions can help you do that. So, John, right, wrong or indifferent, I'm old enough that I worked in an organization that was a little lagging in technology early in my career. And that organization swore up and down that faxes would never take the place of legally signed documents that were couriered. So they were resistant and accepting any legal documents that that were not, you know, hand signed, indifferentiated that were then delivered via Courier. And I'm sure your old company loves DocuSign then, right? Oh, yeah, so, so it's kind of funny, you know, you start looking at all E signing, you look at, you know, just that as an example.

These technologies will find a way. They always do. Mobile phones found a way. Everything always finds a way. So there's absolutely ways that we can do this in a secure manner, that you can put governance around that just like any other application, and you can use it in a very positive manner and in a way that really makes users more productive, whether that be organizationally, specific use cases, but it really does have that value. Now, one thing I do want to bring up when we talk about generative AI and we're monitoring this generative AI, what Microsoft and other companies have done is really look at at models.

So there's a discussion out there about what is the right or appropriate model to use. And that really depends, that really depends on what you're trying to accomplish with it. Generative AI1 tool does not necessarily do everything. However, conceptually, there are ways that you make generative AI more relevant. Open AI and ChatGPT really kind of drove this revolution because they created a model that was bigger than any of everybody else's model by huge factors. So it went from your generative A model before to open AI's model, right? It just grew exponentially very quickly, and you have the ability to reason over that data very, very quickly, well beyond just model, right?

So they're different models that do different things. Purpose built models may be more appropriate if you have something that's very specifically addressed within that use case. But understand there are ways that you take large models and make them more relevant through concepts like building skills. Everybody remember the concept. A lot of people know it from the Alexa, Amazon Alexa. And building skills. Well, you have the ability to build skills, you have the ability to ground data. There are ways to take these very large generative AI models and make them more relevant to the individual or the use case.

You can start taking information like I'm going to talk about APT. I mean, from a cybersecurity perspective, Advanced persistent threat, throw out anything that's talking about APT is the abbreviation for apartment and make that information more, more relevant to the use case that's in there. So I would say, while we do have discussions around model and it really comes down to what are you trying to do with it and how does it work and are you training it and where does that that data go? You're also looking at applications that run on top of it because these applications, Microsoft's copilots run on top of open AI's model.

It's the same model that ChatGPT uses. We're putting it into a SAS application. We're building these applications that make that data more relevant. Well, there's a whole realm of assessment around appropriate AI tool, which is not just appropriate AI model, although there's a whole nother discussion we can have about small language models and pushing models to the edge, pushing them on devices. Running an AI locally on my phone, it really is about where that information is coming from, eliminating bias and then how is that information being disseminated within the the context of the tool. So I think transparency is really the answer to a lot of that.

Transparency and how users are using tools is where AI outcomes comes in. And then you use your reports and your other aspects of SoC 2 reports, things like that to determine it's compliant. And you can merge those together into an AI policy that will benefit both the users, protect users rights, protect the employees, protect the organization from data leakage, and still allow these tools to be used in a very productive manner. All right. So, so if I could just ask one question for you, John, are you starting to see questions on data loss associated with any sort of in house generative AI applications?

Are you starting to see organizations having that discussion yet? Yeah, I think this goes back to how organizations are looking at their overall AI, you know, strategy and and policies, right. So, you know, naturally there are some organizations just because what they do, how they're regulated, who they do work for, they may look at, you know, on premise AI solutions. I think this goes back to you. There's a reason why we like SAS solutions, right? The the rapid acceleration technology, the flexibility, the adaptability of them. So we do see some organizations looking at those.

I would imagine, you know, organizations who, you know, do work in the aerospace and the fence arena, maybe those have, you know, defars regulations looking at, you know, on premise AI solutions. I go back to my look at the vast majority of organizations, you know, there's a reason why they like SAS solutions right there. They're they're quick, they they rapidly, you know, are able to to flex to their needs, they're accessible. And so I think the vast majority of organizations are actually looking more at, you know, those publicly accessible solutions and saying what's the right licensing model to to protect our data and keep our data residency overall, Mike, I guess in the conversation with your clients, you see many of them looking at, you know, on premise AI solutions.

We do, right. So from that perspective, what will happen is either they'll build because of that residency. Microsoft does, for example, have an AUS government version called Azure Open AI for Gov, right? So it's AUS government IO4 approved Ato product inside that realm. But those applications that they're running on them, they now need to start looking at, well, how do we ensure the application that's utilizing this general model is also compliant? So from that perspective, now they're like, well, how do I capture the data? How do I analyze the data?

And it all becomes part of purview as well as there's some some announcements coming at Microsoft at night in November that'll go into some of that other capability. But definitely the the risk of sending data and home grown apps has has started to get started to get a lot of momentum as individuals want to purpose build things for their very specific use cases but still remain compliant and remain below their risk threshold. All right. So maybe in summary of what we talked about today, generative AI, what is it? It's a type of AI that can create new content from user prompts.

It supports a wide variety of use cases across business functions. It's increasing in usage both in business and in day-to-day life, and it presents risks related to data security and compliance issues with Purview. AI Hub, Per View provides tools to enable information protection, data governance, risk management, and compliance. AI Hub provides visibility into who is using what AI tools across your organization. AI Hub can also enforce rules to limit access or functionality of AI data of AI tools to protect data and it can enable you to support compliance management with controls to handle data according to its regulatory.

Requirements. So as we move into questions, John and Mike, are there any closing remarks you'd like to make or you can start looking at some of the things including the QA? Yeah, maybe I'll just kind of start as a recap. I love Mike's example of how the organization that says, you know, you know, back signatures would never, never take care of take care of hard copies. I think, you know, the key item here is, you know, Gen. AI solutions are are here to stay. And we're probably only, you know, barely scratching the surface, right.

I, I personally have trouble remembering what it's like to do my job without these solutions at this point, right. So they're here to stay is the more rapidly you figure out an overall strategy for the usage at your organization is the better. Naturally grounding this in policy is key. And I think from there, and it's the identifying methods where you can make sure that you're you're conforming that policy too, because with great the great capabilities always comes great risk, right? And that's with everything, right, whether it's SAS solutions, leveraging Jenny I capabilities or anything else. So we encourage everyone, you know, we wish we did have that point question, you know, how many organizations have that kind of policy in place today?

Because we do see organizations struggling, right? So, you know, the sooner you get ahead of it, probably the better because organizations, you know, everyone in your organization is probably using it for something as we speak. You know, it's interesting, is a Microsoft employee and a tech security technology professional. For a long time I have often focused on the technology tools, right, the bells and the whistles, the features, the things that are going on. But I think more than anything, what we, what I've really experienced is that organizations are having to be a little bit more thoughtful about having discussions about what should that policy look like?

What do we want to do? And for me, you know, where I get brought in is, hey, how do we dispel some of the myths and some of the fears? What keeps the legal counsel up at night? What keeps the SISSO up at night? And then talking about how AI tools align to that so that they can still see the the positives, right? So the employees can use this, they can get benefit. I deal with the naysayers all day every day, you know, trying to explain how these different tools work to to the folks who are paid to worry, right. So ultimately that has been the key.

So I think what you said about having a policy and then configuring the tools to align with that policy has, is really a, a good thing. Getting involved, You know, if you're an international organization or you have California employees, you have European employees, GPR, things like that, Understanding what those employees are needing and then what the organization is needing and putting that into a policy. And then Microsoft, for example, making tools that are flexible enough to reflect back that policy and reduce the risk for the organization. Organization is that first step. So I I think it's a great discussion to see as we continue to see AI make its way through society.

Great. All right, Mike, can we do have a few questions to come in, Maybe I'll I'll start teeing them up to the group. One that came in was, is AI Hub understanding of AI usage limited to only Microsoft Copilot? Mike, would you? Like to take that so AI Hub right now it's between 300 and 400, It's published online. We actually have between 300 and 400 sites that we're generating. When you really get down to it and you talk about the hardware that's running these services, there is only so much of it in the world.

It is a supply chain nightmare as people have read in the news, getting GPU's and mass scale GPU's. So no, it is not. And we have seen integration even in Microsoft. We have AI Hub, we have Defender for Cloud Apps,

which is a Casby that allows you to look at this. We know where these things are, we're tracking them and we're tracking the compliance of those tools based on the information they provide, You know, out to the out to the community. There's a community of organizations to talk, talk together and provide information on their Sock 2 reports, you know, Fedramp status, compliance with Dora, you know, alignment team, that new EUAI standards, things along those lines.

So no, it doesn't. We track all that. We really focus on how people are interacting with information, particularly the browser and and focusing on that information. Great. Another question that came in, I think this is back to the finance use case around preventing folks from uploading or copying, you know data in in the purview. Does a solution allow light only block any upload or copying, or does it allow you to potentially allow you to block only uploading certain sensitive kinds of information? So it's about types of information. So the real key behind this is, you know, we've we've seen this through different technology iterations over over time, right?

If you just talk about virtual desktops and cutting and pasting between computers and applications and endpoint DLP concepts are around for been around for a long time. Microsoft's approach to this is to allow you to be flexible in determining things called sensitive information types, types of information. We also allow you if you're down the journey of labeling content and creating sensitivity labels, that becomes part of the equation. So it's not just about blocking all data that limits the flexibility and productivity of organizations. However, based upon risk, you can scope these rules differently. Different rules can go to different people and that allows you to determine where does your risk sit with certain individuals, certain groups versus versus others and ensuring that everyone is aligned appropriately and you can you can adjust that based on job role, geography, etcetera.

Great. And probably our final question, because we argue about running out of time in a minute or two, can these solutions also help protect uploading of non office or PDF files and then getting granular that astute as well? Yeah. So there's a lot of questions around this. So around when you're dealing with cutting and pasting insensitive information types and interacting with these AI services, that's not just about files, right? That's about the content that's moving there. And the browser extension is going to catch that, is going to take care of that.

There's a, we can have some pretty detailed technical discussions around that. So trying to keep it white. There is however, the ability to interact with files. And if you're taking files out of a labeled file, so taking it out of Word and putting in that, well, now it's looking at the sensitivity label associated with that, with those protections that are on those types of files. There's a whole lot more information about what Microsoft is doing for non office files. PDF has been integrated fully. So all that is available in PDF.

But when we're talking about other files, we're we're going down the road there. Ignite in November is where you're going to want to hear all that information. I can have NDA level discussions with customers, but Ignite just know that if content is coming out of a sensitive file, that file could be analyzed if that is your choice or you can look at that information. So now we're not looking at at the actual classification of the document, looking at what is the, what is the information here and your sensitive information types can reflect into that. So now we could say, is that a Social Security number?

Is that whatever it may be, whatever keeps you up at night, it can look at that and make that determination separately from metadata associated with the file. Great, appreciate that. Well, it looks like that's last question and then we are out of time. So we'd like to thank everyone for joining us today. If you have any following questions for myself, Jose or Mike, please feel free to reach out at any point. This session is recorded will be available online as well. And thank you for joining us today. That concludes today's webcast. If you have questions or would like more information after the conclusion of this webcast, please feel free to contact any of the presenters using the information found in the resource widget.

We want to thank the audience for attending and for participating in today's event. We would also like to thank today's speakers for making today's interactive webcast possible. Concluding today's live event, you'll be asked to complete a short survey. We encourage all feedback and suggestions to better improve our webcast experience. The link to download the slides remains available for a limited time, so obtain your copy today. We will also share the recording and slides via e-mail. Be sure to connect with us on our site or follow us on social media to stay up to date on future events.

We look forward to seeing you on a future Grant Thornton webcast.