[Kevin Szczepanski]: Welcome back, everyone. Well, we are excited—I am excited to have Nick Reese with us today. Nick is, in a phrase, highly decorated. He will not like my saying that, but I'm going to say it, to all of you. And you'll understand when I finish this introduction: for ten years, Nick was an operations specialist for the United States Navy. He was a member of the US intelligence community for several years, including a stint as director for emerging technology policy for the US Department of Homeland Security. He is now a board member of the Homeland Security Advisory Board. And maybe not most importantly, but arguably most importantly today, he is co-founder and COO of Frontier Foundry. And we're going to ask Nick to explain a little bit about what Frontier Foundry is. But welcome to *Cyber Sip*, Nick.

[Nick Reese]: Thanks, Kevin. It's a pleasure to be here.

[Kevin]: It's great to have you. Let's start with a little bit about you. If you were in an elevator and had to walk me through your background, tell me a little bit about your background and tell me about Frontier Foundry...

[Nick]: Sure. Well, if it's the elevator speech, then I'll have to say it's always been about pursuing things that I care about it to and having the most impact on those things that I can. And so that originally started in the military. It went into the intelligence community, and then eventually I moved into the policymaking community. And it was there that I kind of looked back on my career and said, you know, the thing I want to do next is build. And so I paired up with my, my co-founder, Sultan Meghji, and we set out to build an AI company because I had been building AI policy, within the Department of Homeland Security, advising the White House, things like that, on that and other emerging technologies for many years. And very proud of that work, very proud of my team. But the thing that we really wanted to do is create something. And so at Frontier Foundry, we set out to build AI that was privacy-preserving. And I mean that in a real way, like really measurably privacy-preserving and the reason we did that is because one of the tragedies of AI development, at least right now, is that the people that deal with the most sensitive data are often the people who are not allowed to use the best tools. And I was one of those people. I was one of those people in the intelligence community. And so we went out and we built, privacy-preserving AI products. We've got a few different ones, including an LLM, and we also built it to be able to run in different computing environments. So you don't need a cloud that's powered by Three Mile Island in order to run our AI, you can run it on a laptop, you can run it on a server, a small server, anything that that you need. And that's about making things accessible to people. And it's about actually addressing mission need, in a real way.

[Kevin]: That's and that's critically important today because as many of us know, and as we're all learning, power equals cost. And one of the real challenges, I think, with AI is making it accessible across the board. It's not just the big companies, the big law firms who have access to it. We all need access to it in order to really lift all of our boats up as we approach the next decade of technology and economic growth.

Season 4, Episode 6: "Yes, You *Can*—But *Should* You? Strategies for Implementing Generative AI"
04.16.25 | *barclaydamon.com*

BARCLAY DAMON LLP

[Nick]: Yeah, that's absolutely right. And, you know, I like that you talked about kind of the… broadening the access to AI. And I want to I want to dive in on something here because I think it's really important. There's democratizing AI and then there's providing access to AI. And those two things are different because democratizing AI is the ability for someone who's never put their fingers on an AI, you know, machine before to use it and to learn about it and to understand what AI really does. And if you look at some of the kind of big box AI models they've done that, right? They've created these public platforms, they have free versions, they have paid versions where you're able to go on and just have a conversation or ask it for specific things, and that democratized AI. But I think access to AI means something else and what I, what I think when I think about access to AI is it's something that is specific to what you actually need. And so it's not this kind of big, generalized model that is fun to play with and can be useful in some instances, but it's about having access to AI that is specific to your needs. And that's, I think, a difference. And I think that that actually represents one of the shifts that we're seeing in the AI market now.

[Kevin]: And it's a critical shift. And let's keep that shift in mind as we move through our conversation, because I think there'll be several points where it is very relevant. So with that in mind, I want to talk to you about the difference between what we call generative artificial intelligence or generative AI, and somewhat similar sounding, but very different, artificial general intelligence. Can you walk us through the difference and, Nick, maybe touch on a little bit why everyone seems to be so fascinated or frightened by artificial general intelligence.

[Nick]: Sure. So this is a really important point. And I'm actually an adjunct professor at New York University and I teach my students this because it's such an important kind of, thing to understand about AI. So AI as a concept. I mean, goes back to Alan Turing. And there have been various different types of AI, different types of iterations of AI for all of those decades. And we've seen different types of…

[Kevin]: …Sorry to cut you off just for the benefit of our audience, Alan Turing, meaning AI in some sense goes back to World War II.

[Nick]: Yeah, absolutely it does. And that's where the initial theories came in and the initial testing. And a lot of what we see today was built on those initial theories and those initial, you know, creations. And so we've seen several different iterations. And even today, AI is not one thing. And it's really important to kind of understand that AI is not a single thing that you can say, well, that's definitely AI. One of the things that was difficult for us when I was at the Department of Homeland Security was working on what a definition of AI actually was. And when you think about it, that's harder than one might think initially to, to really define that. But when we think about artificial intelligence, we think about intelligence like, you know, a machine that is intelligent in the way that you are intelligent.

[Kevin]: Yes.

[Nick]: And right now, that's not what our AI is. Our AI thinks differently than us. It is better at things, certain things. But then again, it's worse at certain things. So it's very, very good at, you know, big mountains of data and making sense of that and providing analytics. But it's not as good at context and things like that that humans are good at. So really, a lot of the benefit when it comes to AI is when you pair those two things together. Now, generative AI is something that we all kind of remember when it came out and, you know, or at least when ChatGPT came out, which was definitely when it became most known. That's an artificial intelligence that actually creates or generates something entirely new. So it's not that it's just kind of reading some data and then visualizing it for you. It's actually creating something entirely different that didn't exist before, whether that's a picture or text or computer code or video or whatever. It does…and that is its function. Now, artificial general intelligence goes back to this idea of intelligence in the same way that you're intelligent or I'm intelligent. And so artificial general intelligence has been, you know, dramatized in Hollywood for a very long time. It's the idea, it's, you know, it's the Skynet and it's the idea that there's this being that is

BARCLAY DAMON LLP

much more powerful and knowledgeable than humans, and it's something that's being pursued. But we have to make a distinction that those two things are very different. And just because we have generative AI doesn't mean we have AGI.

[Kevin]: Right? So when I'm interacting with ChatGPT or a private service, I'm interacting with generative AI, artificial intelligence, general artificial intelligence. I misstated that. Artificial general intelligence is probably—depending on who you talk to decades away, maybe more. So what's your take on why, when people use AI, they're using generative AI. But it seems to me in many cases when we debate AI in the public square, we're debating the Matrix or the artificial general intelligence that everybody's worried will take over the world, supplant human beings and cause everyone to die. What's your... you're teaching students... What do you hear from them, and what's your take on why that seems to be the case?

[Nick]: Well, I think one is just a gap in education. And so I think there's just a... and a fundamental understanding that we need to be able to get to. And unfortunately, we've had a lot of Hollywood movies and series and different things like that that have portrayed very good stories that have the concept of "artificial general intelligence," as its core. And so there are a lot of people for whom that is their touchstone for artificial intelligence, right? They don't really make the distinction that artificial intelligence and artificial general intelligence are different things. And so I ...you know, I think my first thing is just making sure that we have the ability to kind of educate people on what this is and what it isn't. And I don't mean that we need to teach people to be algorithmic engineers or computer scientists, and they don't have... I'm not saying they need to be able to code it, but I do think that they need to understand what it is and what it isn't. And so I think I really like the point that you brought up about, you know, if generative AI is what we have, then why are we talking about it like it's AGI when we don't have that yet. And I think that's a, that's a really important thing that I hope that listeners really think about from this podcast.

[Kevin]: So let's leave AGI aside for a moment, because we don't have that today. We do have generative AI. Let's talk about that. And I think one of the first questions I want to ask you is, how does an organization—and let's focus on if we can focus on small to medium sized businesses—I've heard from so many clients an urgency about adopting AI. "We need to do this because we need to stay ahead of the curve." "We need to avoid falling behind." But that doesn't strike me as the right way to start thinking about AI or AGI. So. Or generative AI. So help us out, Nick. If you're a small to medium sized business, you're talking to them every day. How do you start that conversation?

[Nick]: Yeah, and you're right. I do talk to these folks every day. And I see kind of two categories of, of I'll say leaders of those organizations. So the first leader is kind of what you described where someone would say, well, we need to do it to get ahead of the curve. And I've actually been hearing this even since I was at DHS, where I would hear people say we need AI. And I would say, okay, well, what do we need? And they would like, you know, you just you just bring some AI and, you know, you sprinkle it on top like it's your seasoning, right? Whatever you're grilling in your backyard. And it's like the people that view it as "I just need it," are the people that have the hardest time. The people that sit back and say, what is the actual use case/business value case that I can make? And one of the things I tell people all the time in my industry role, and in my academic role, is the technology should not dictate the vision. So if you have a vision for your organization, so whether that's getting more value out of the mountain of data that you're collecting or it's doing something faster, or it's streamlining a process, you shouldn't allow the technology to say, oh, actually, you can't do that. And we see this a lot. And there's a I think a big thing, especially in industry, is there's so many people that have this like artificial barrier to entry in front of them that they think is there but is not necessarily there. And I think it goes back to the way that we've been building AI, especially generative AI, for the last few years.

[Kevin]: Tease that out a little bit more for us, Nick. What is it that actually exists as a barrier, or what is the perceived barrier that small to medium sized businesses face?

Season 4, Episode 6: "Yes, You *Can*—But *Should* You? Strategies for Implementing Generative AI"
04.16.25 | *barclaydamon.com*

BARCLAY DAMON LLP

[Nick]: So first of all, data security and data privacy, that's no longer kind of a "nice to have." That is something that is core to your reputation, to the function of your business. You have to have those things. And so whether you have regulated data, like financial data or HIPAA data, health care data, or maybe you have, attorney-client data, like in a law firm, you know, if you have those or even if you just have data that's, you know, a little more bland, some customer data—protecting that data is something that all organizations are looking at and saying, we have to do this now. And so a lot of organizations are looking at generative AI and they're saying, well, some of these bigger companies will sell me a version of their large language model, you know, publicly available large language models, like, you know, Claude or ChatGPT or something like that. And they'll install that. And the problem becomes, well, where does the data go? So, so I use it. So I take my data and I put it in there...where does the data go? And that question is a huge barrier to entry for a lot of people where they say, if I can't say exactly where my data is going, then I can't use this product. But I think the problem is maybe you can't use that product, right? But you can use generative AI. It just the idea that using generative AI is equal to a loss of privacy is incorrect. And that's one of the points that I spend a lot of time talking to small and medium sized businesses about.

[Kevin]: I see what you're saying. So if I'm a small to medium sized business, I might be coming in with a perception that generative AI is ChatGPT or some similar public-facing AI platform. And if that's the case, I might very well conclude—correctly—that I'm unable to use this platform because I don't know where this data goes. And if someone, including a regulator or plaintiff's lawyer, asks me what I've done to protect the data, I'm not going to have a good answer. Right? The best answer I might have is, well, it's in the cloud, so someone else is protecting it, which we both know is never the right answer. But that doesn't mean there isn't a private generative AI platform that can be used, and it's closed in the sense that your data is protected.

[Nick]: Yes. And this is boring, but I think it warrants mentioning which is that really it's about how your system is architected. And I know a lot of people are probably already their eyes are glazing over. But if you can imagine, you know, your data starting at your organization. So let's say that you have a secure cloud that your organization pays for. Your... you have that cloud infrastructure, that is your cloud infrastructure only, and it has protections around it that you dictate. Now, if you use some of these cloud-based, cloud native, artificial intelligence products, that means that you're going to leave that environment and you're going to go to a like a different cloud. And then, where it gets really questionable is, then once that data gets to that cloud, is that data then going on somewhere else to be part of a training regimen for another algorithm? That's a really important question. And it's actually an architecture question, which again is boring, but it warrants I mean, the way I would say it is: how many different fingers touch your data? Is it just yours? If it's just yours, then I think you're on the right track. If it's if you can count, you know, on hopefully just on one hand, hopefully not on two hands, the number of hops that your data is making then, then you really do have an issue.

[Kevin]: So let's not role play but let me use a law firm as an example. So let's say I come to you, and I say, all right, Nick, we're in the process of evaluating AI products. And it's not... we're not using ChatGPT. And we know that there are AI options available on other platforms like Lexis and Westlaw and others. But we're really trying to identify that use case. And let's say that so far we've looked, and we really haven't found an AI platform that gets us any further than where we are right now. It can replicate what we do, but it doesn't do it faster enough. It doesn't save enough in costs, and it doesn't secure our data enough that we're comfortable using it. So that's where we are. And I come to you and I say, all right, so how should I be thinking about this? What are the benefits? And I suspect you might say, well, you tell me what the benefits are, right? It's because I have to develop the use case. But if I'm struggling with what are the benefits and weighing those against the risks of AI, can you walk me through that?

[Nick]: Yes. So I think the place I would start without... so like before, the use case identification process, I would start with asking you kind of I mean, admittedly, kind of a rhetorical question, but I think it's important because it illustrates something. So, you know, if you were a law firm that does primarily contracts, let's say that's your main business area, and you needed to have... to hire a human, a lawyer to come in and do contract

Season 4, Episode 6: "Yes, You *Can*—But *Should* You? Strategies for Implementing Generative AI"
04.16.25 | *barclaydamon.com*

BARCLAY DAMON LLP

evaluations for a certain type of contract in a certain state. Would you hire a generalist lawyer, or would you hire a lawyer that is a specialist in contract law in that state? But the answer is you would hire the specialist, of course. And all organizations do this. You hire a specialist over a generalist when a specialized problem presents itself. And so I would ask you to apply that same logic to your AI. So if you're going to bring an AI in, do you need a generalist AI or do you need a specialist AI. And so I would posit that maybe one of the initial reasons that you're not finding what you're looking for is because the AI that you're pursuing is not built for your purpose. And when I say "built for your purpose," I mean not one of these big kind of general trillion parameters trained on the entirety of human knowledge type of AIs, but also these kind of bigger, in your case, legal AIs, where, you know, you have Lexis and Westlaw and some of these other, you know, Relativity and some of these other things. I realize those are legal specific, but that's probably also not good enough. You need something that is not just broad-based legal. You need something that's specific to exactly what you're looking for. And it's possible to have an AI that you can tune to your needs. And that's the thing that I think is the most valuable, when you can kind of say, hey, this generated response wasn't exactly what I was thinking. I think it needs this. Great. Go back. You tune the model, you come back, and now you have a true kind of human in the loop. So now you have a specialized AI. The algorithm is now tuned to a point where it is actually proprietary to you, and you have customer trust and client trust there because you were involved in building it. And so that's the place that I would start that conversation. And then I would probably also talk about, you know, the... making sure that the technology doesn't, dictate your vision.

[Kevin]: Right? I think that's a great distinction. General versus specific. I find it very useful in understanding it myself. So I guess then what the takeaway should be for most law firms or most SMBs looking to incorporate AI into their business, is you not only need to have the use case. So the use case is "I need to review and compare hundreds, if not thousands of different contracts to look for trends and similarities and differences," but then past the use case, I need to understand how that system is going to work, and my understanding of how it works is critical to implementation and security.

[Nick]: Yeah. So one of the things that I, again, I like to talk about is, you know, there are people in the world that are car people and they're people who are not. I am one of those people who is not. I can drive a car. I can generally operate one, but I am not the person you want popping the hood. But I know enough to know when I need to take it to a mechanic. How do I know that, if I don't really know much about my car? Well, because I have a bunch of gauges in front of me, right? And those gauges tell me different things. And so if it's overheating or the tire's low or whatever it is, I know that, oh, I need to go seek someone to help me before this becomes a huge problem. That doesn't really exist in a lot of the big box AI companies, you can't look at anything and say, I think I'm getting a biased answer, or I think this is hallucination or something like that. There's nothing you can look at. And so this is another thing that that, you know, we were kind of puzzled by when we were building our company. And so another thing that I would really recommend that people consider is: do you have the ability to audit your system. So do you know what data was trained and when? And if you don't know that, you have a hard time tracing back, hey, I think this was a bad response. Or hey, I think this is a hallucination or something like that. Do you have the ability to do that? You know, do you have the ability to bring up the sources and bring up the documents and like, see the source material with your eyes. And you know, when you're in some of these—whether it's regulated industry or, you know, an industry like, like legal, you know, very ethically bound, you have to ...that's not an option. We had a customer that was a hedge fund. And if the SEC came calling, they couldn't just say, well, you know, the black box said so. So we made the trade like, that's not good enough, right? You have to have more than that. And so that auditability function is also what builds trust with your clients. But it also creates... it is a reputational boost to you.

[Kevin]: So this leads to something that I try to tell everyone that I see, including clients that we talked to about AI. You need to understand what you do, whether AI can improve it. You've talked about then the next step, which is all right, how specifically will I work and what data are you inputting into this system? And how are you protecting it? Now, you mentioned hallucinations and bias, and I think many people know maybe not everyone, but many people know that one of the potential downsides to an open AI system is the potential for

Season 4, Episode 6: "Yes, You *Can*—But *Should* You? Strategies for Implementing Generative AI"
04.16.25 | barclaydamon.com

BARCLAY DAMON LLP

fakes, which we call hallucinations, bias, which essentially means the old phrase garbage in, garbage out. If your data reflects a bias, then the responses you get may reflect a bias, and you may not know what that is if you're not careful. That said, Nick, can a private AI platform protect against bias or hallucinations better than a public platform like ChatGPT?

[Nick]: Absolutely. So let's talk about hallucination first. So one of the reasons that you get hallucinations is because you have such huge training data sets that are not related to what you're asking it for. So again, let's go back to our law firm with the contracts in a certain state. If that's what your use case is. And you need to do those comparisons to find the differentiations and all these, you know, that's what your use case is. Do you really need to train your eye on the entirety of X or Reddit or something like that? Probably not. And so when you when you bring in all these disparate sources, you sort of... you start to lose control of the types of responses that you probably can expect. And so this is where you run into problems with hallucination. So what you can do is you can reduce that risk considerably by understanding what your training data set is and by being able to audit it. So, you know, like our large language model was designed specifically to have input into it, take, you know, your data and put into it and create that specialist. Right. But it also... it is trained to still speak English. It understands tone, it understands sentiment. But it's still important that you know where that came from. And so that hallucination mitigation is of huge importance to any user of artificial intelligence. So then on the bias side of things, this is this one's a little more difficult sometimes because you... the AI is not malfunctioning when you see when bias occurs,

[Kevin]: ...it's that you're getting it exactly as the data suggests, right. And so data is biased.

[Nick]: Right. And so... and it's pulling things out that you didn't see because you're human and you don't think the same way. So this is where the importance of the human in the loop is, and this is what—when I was talking about like the tuning process where, you know, on a first iteration, I would bring you the first draft of a generated document and I would say, all right, you know, Kevin, this does this look right? And you go, wait, something's weird about this. Then we can go back under the hood and we can figure out what happened, and we mitigate that before it becomes a problem. And so we still don't know that there will be perfectly unbiased data, right? There may still be outputs, but that's where the importance of the human in the loop is. And that's also where the importance of understanding what your training data set is. And that's I think, you know, we've heard the story about like hiring and things like that where they run into problems. And I think understanding what some of the historic issues with that data might be is a great way to also make sure that you're mitigating that risk.

[Kevin]: It also makes me hark back to what we were talking about earlier, the distinction between generative AI and artificial general intelligence, the notion that you can implement generative AI to execute your use case, make your business more efficient, cost effective without human oversight is a fallacy. You have to have that oversight in place. So let's turn to that if we can. And I want to expand the discussion to what we call AI governance or corporate governance. Talk to me about that. How does a business get its arms around its employees, its use of AI, making sure that the actual implementation is consistent with best practices.

[Nick]: Starts with asking the right question and the right question when it comes to an AI system is "should I" not "can I." And so let me unpack that. So what I mean by that is... let's talk about IT procurement, you know, phones, laptops things like that. The way that IT procurement works is you know, there's a set standard of types of encryption or logins or hard drive size or all these things. And if that piece of equipment meets those criteria, then the answer is yes, right. It's a "can I" question, but there are some organizations that look at AI the same way where it's like, oh, does it have certain security parameters? Does it have, you know, what's the size of it? What's the cloud requirement, GPU requirement, all that stuff matters, don't get me wrong. But if we're talking about governance, we need to be asking, should I use AI for this use case? And if yes, under what conditions? And governance is a very personal thing to each organization, because what it means to a law firm is not going to be the same as what it means to a, I don't know, a store or a bank. Right. Those are

Season 4, Episode 6: "Yes, You *Can*—But *Should* You? Strategies for Implementing Generative AI"
04.16.25 | barclaydamon.com

BARCLAY DAMON LLP

going to be very different. So you can't really prescribe this is exactly how you do it. But I think starting with the right question is the first step. And then what you do is you need a way to kind of frame the discussion in terms of how much automation versus how much impact does your system have. So is it highly automated and highly impactful? An impact means different things to different organizations, right.

[Kevin]: I was just going to ask you about that.

[Nick]: Yeah, it could be. It could mean loss of life depending on where you are. It could mean financial loss. It depends on your industry. But if you have high automation, high impact, that system probably needs extra governance. If it's low automation, medium impact, then that means that... by definition means there are humans involved anyway. And so maybe the governance is less on that. And so you need to be able to measure the impact versus the level of autonomy of a system. And then ask yourself that question should and the conditions would come in where you say, okay, this is a highly automated system, so we're going to use it, but we are going to perform data audits every two weeks, something like that.

[Kevin]: When you... And that makes sense Nick. And not to put you on the spot. I was thinking when you were saying that, that high automation is something that many health care providers are increasingly interested in, but there's an example of potentially high impact. You could be talking about life. You could be talking about disease versus cure. I'm curious. And maybe you haven't, but if you've encountered either health care or the health care industry considering AI or some similarly regulated industry, and how does the evaluation of the use case differ between, say, a highly regulated industry and perhaps a manufacturing industry where it's important, but they may not be subject to the same regulations?

[Nick]: Sure. So less and less. I have less experience directly with health care, but I do have experience in regulated industries, for sure. So we actually work quite a bit with... in a regulated financial services industry. And so subject to SEC regulations and things like that. And so I, you know, I hate to go back to this, but I mean, honestly, the reason that that project was successful is because the leadership in that organization had a vision. And they said, you know, this is what we want to create. This is what we have now. This is what we want to have in the future. And they were willing to say, here's the data. And then here are parameters around that data. You know, you can't do certain things with it, right? They gave us the parameters. But then we built it. And, you know, it took way less time than people think. And there was, you know, the leaders, and the people in the organization were involved from the beginning. And I think that kind of partnership, when you create that, that is where you end up with something that is not only more effective, but it's more easily governable because you built it kind of together. Where I think you run into trouble is where ...is what we're used to when it comes to software as a service products, where you kind of, you know, you walk up to the shelf, you put the product in your basket and then you check out. Right. And the product is what the product is. In some cases that's just fine. In some cases it's not fine. And I think particularly for organizations that have sensitive data, regulated data, or I mean, just are really concerned about their reputation. You know, you can't just say, you know, buy the thing off the shelf and then get a governance structure off the shelf to go with it. That's not going to work, and it's not going to be as effective, and you have a higher risk of running into trouble, kind of externally.

[Kevin]: Couple of quick questions, Nick, are you seeing the businesses developing AI committees, implementing AI policies? Is that a good idea for any business that's even considering purchasing or implementing an AI product?

[Nick]: So yes, I've seen it. Actually, specifically in the legal industry, I've seen a couple of firms that have these committees, and I think they're a good thing. And I think it's a good thing because organizationally, there's a couple things you need to establish for yourself. And so one is kind of like your use ethics. So what are kind of the things you're going to use AI for? The things you're not, the types of AI that you're, or the types of data you're going to use, types of data you're not. And there's no like right answer to that. It's just you have

Season 4, Episode 6: "Yes, You *Can*—But *Should* You? Strategies for Implementing Generative AI"
04.16.25 | barclaydamon.com

BARCLAY DAMON LLP

to define it. Right. And then that organization, that committee usually becomes the governance committee. So you've already got the people in place to create the governance. And so I do think these are good things. The technology obviously moves fast. So I always kind of recommend that people consider bringing in outside experts, academic experts, industry experts, something like that, just to make sure that they're up on the technology, the policy, executive orders as they come out, things like that. But I do think it's a good idea, and I think that helps create that vision, that then you go out and you look at the market and you say, okay, which products help me execute this vision?

[Kevin]: Yeah. Makes sense. All right. So Nick, we are almost out of time. Your knowledge is so vast. I feel like we've only scratched the surface. So let me give you one last opportunity and I'll keep it very general. Sort of. You're coming in. You're asked to come in to speak to an SMB about AI, and they know they want to explore it, but... and they may have an idea about the use case, but they're very cautious. They want to make sure they know the strengths and weaknesses and how they should be going about implementing AI strategically. What's your opening address to the board of that company sound like?

[Nick]: Start with a pilot. Start out with something small scale, but be willing to go through the process to build it in partnership with someone that shares your vision. And if what you're interested in is expedience, just getting AI in as quickly as you can, that's probably a recipe for failure or at minimum setback. But if you're willing to put in the time to define the use case, you're willing to, you know, not allow the technology to dictate your vision and you're willing to kind of find a partner that agrees with that vision, that can deliver on that vision. Then you start with a pilot, prove it out, prove it out. And it's okay to say, you know, we're going to run a two-month pilot for way cheaper than buying an enterprise solution. None of that work is lost. You will have learned a lot throughout it. You can even build your governance as you're building the pilot. So you even have like governance pilots, so to speak. But I would say to start small and to make sure that you are not... you do not believe that there are artificial barriers to entry in front of you. Be open to what the kind of new state of the technology is, and be open to the idea of specializing over generalizing.

[Kevin]: Well, that is smart and sage advice. Let's leave it there. Nick Reese, Frontier Foundry, thank you so much for joining us. I really appreciate your stopping in to *Cyber Sip* today.

[Nick]: It's a pleasure. Thanks for having me.

[Kevin]: And we'll continue this conversation. Thanks to all of you for joining us in this episode. We're back soon with another one.

[Kevin]: The *Cyber Sip* podcast is available on barclaydamon.com, YouTube, LinkedIn, Apple Podcasts, and Spotify. Like, follow, share, and continue to listen.

This material is for informational purposes only and does not constitute legal advice or legal opinion. No attorney-client relationship has been established or implied.

Barclay Damon Live *podcast transcripts and captions are automatically generated through artificial intelligence, and the texts may not have been thoroughly reviewed. The authoritative record of* Barclay Damon Live *programming is the audio file. Thanks for listening.*

Season 4, Episode 6: "Yes, You *Can*—But *Should* You? Strategies for Implementing Generative AI"
04.16.25 | barclaydamon.com

BARCLAY DAMON LLP