

Barefoot Innovation Podcast: Sean Neville, Co-founder and CEO of Catena Labs

***Note that transcripts may sometimes contain errors and that transcript timing notations do not match the posted podcast**

- Jo Ann Barefoot: [00:03](#) I have been looking forward to today's show for a long time because we have a really exciting guest. It is Sean Neville, who is the founder and CEO of Catena Labs. Sean, welcome to the show.
- Sean Neville: [00:18](#) Thank you so much. Thanks for having me.
- Jo Ann Barefoot: [00:19](#) I'm really thrilled to have you. You and I have known each other for a long time. I lived in Boston for two years when I was at Harvard working on a senior fellowship and I lived in Seaport, and I was just a few blocks away from your charming old office at Circle and we used to run into each other in the neighborhood and I know Jeremy and John Beccia and the original team. And what you have done, and I was an early customer of Circle also, so I got a first-hand taste of better payments might look like back when you had that original payments model. So I think a good place to start, so you were the co-founder of Circle with Jeremy Allaire, so let's maybe start with that a little bit. Talk to me about what was the problem that you and Jeremy and your team were trying to solve at Circle and how has the vision for USDC evolved and where are we today?
- Sean Neville: [01:34](#) Yeah, so my background is in software engineering and Jeremy and I have known, my co-founder Jeremy Allaire and I have known each other for many years, but we both became more or less obsessed about the same problem space in 2013, which is when we started Circle. And the idea was that we had spent the early part of our careers building on open internet rails, allowing people to distribute content, expressing their opinion, sending and accessing data and so on and so forth over open internet rails. And we wanted the financial system to function in a similar way and had imagined what the world might be like if everyone around the world could tap into a global financial system that was effectively borderless and nearly free just as emails are borderless and nearly free to send. And what would that mean to individuals, but also to businesses?
- [02:34](#) And we sort of imagined that as a new kind of prosperity if such a thing were possible. And at the time I wouldn't say we knew exactly how to make that happen. We were just committed to

the space. It wasn't until a few years later that we developed US dollar coin, USDC, the stablecoin. Maybe it should have been more obvious to us in 2013 that a digital dollar what was needed, but we did start with Bitcoin and Bitcoin running on global blockchain rails at the time. Vision hasn't really changed to be honest with you, Jo Ann. It's the same vision that we started with, but the course to achieve that reality has had a lot of twists and turns over the years.

Jo Ann Barefoot: [03:19](#)

Yeah. And I'll just say congratulations on the spectacular success of the Circle IPO, which has really been remarkable this year and this spring. Really remarkable and I think for many people, feeling like stable coins are coming into their own, that this is a breakthrough type of moment. So anyway, it's been very exciting to see. So you left Circle, their day-to-day team in 2020, and did you already know by then what you were going to do next? And could you see perhaps better than the rest of us were where we were headed with AI? Talk a little bit about your logic in founding Catena Labs and what you're trying to do.

Sean Neville: [04:14](#)

So in 2020, I moved into a board role at Circle and out of a direct management operating role. And it really gave me space to explore what was next in financial infrastructure, but also in things like identity primitives, which were very relevant to what we're working with when it comes to AI agents today and wanted room to-

Jo Ann Barefoot: [04:37](#)

Can you explain identity primitives to our audience?

Sean Neville: [04:39](#)

Certainly, yes. And it's effectively this notion that if I want to be sure that I'm talking to you and you want to be sure that you're talking to me and we can't see one another, we're interacting by sending messages over the internet, how can I verify that it is actually you and that no one is in the middle of our conversation? No one's pretending to be me, no one's pretending to be you. How do we handle that on the internet? And there's no standard protocol for that on the internet. If we access a website, we go to Amazon.com, there are ways for us to be sure that this is Amazon, this is really run by Amazon. E-commerce runs on SSL, we have these certificates, but there isn't really an identity layer for human beings or businesses on the internet. It's been added on top and approached in different ways.

[05:32](#)

And so this idea of how do we prove who we say we are to other people, how do we become trustworthy and how do people verify that is just a fascinating space. It has always been a fascinating problem to me. So I wanted to dig deeper into

exploring that as well as what was next for these stablecoin financial rails that we had helped initiate at Circle. Built USDC in 2017. What was the next stage for really making stablecoins useful, not just for crypto capital markets, we may have been very successful in that use case, but what was the next step to get them to true payment use cases and other use cases, programmable money use cases that had excited us about creating Circle in the first place? How do we get there? And I wanted to take some time to think along those lines, and ultimately that landed us with the company that we're building now.

Jo Ann Barefoot: [06:28](#)

And what does it do?

Sean Neville: [06:30](#)

So Catena Labs is building the first fully regulated AI-native bank. And what that means is that we're working on financial infrastructure that is designed from the ground up for autonomous and semi-autonomous digital actors and the people and businesses that deploy them. So we've become familiar with AI. There's been an explosion of interest in large language model workflows. But as these AI systems become participants in the global economy, they begin to need to custody assets, make payments, negotiate payments and handle funds flows. How do they do that safely and reliably? How do we prevent bad actors that are leveraging AI from accessing our systems? How does this work? And we're building that infrastructure from the ground up, and it's a very exciting space to play in. Certainly stablecoins have a role. It turns out stablecoins are quite useful as machine-native money. So if we have machine intelligence, it makes sense that they're not necessarily handing over a credit card, but they're leveraging things like digital dollars on the global internet rails. So there's a nice synergy there as well.

Jo Ann Barefoot: [07:48](#)

So it seems to me that you, more than most of the people I know, maybe any of the people I know, you are right at the nexus and the cutting edge of two of the biggest things that are changing our lives in technology. You're working with stablecoins and you're working with AI agents, and either of those is fascinating and important, but when you put them together, you're trying to create a truly different way for the system to work. Can you go deeper for us and talk about what's that going to look like in our lives?

Sean Neville: [08:36](#)

Yeah, so it is true we're operating in between two sort of surging demand curves. One force that we're working with is that the dollar is going internet-native. And now we have clarity in terms of United States regulation thanks to the GENIUS Act which

recently became law. We have clarity around what exactly is a digital dollar on open internet rails. And so the dollar becoming internet-native is one force, but the other force is the internet itself is becoming agent-native. And we're in this phase where the web is really shifting into a new kind of era and it's affecting everything from the tech and the way we interact with one another and content to the business models that run the web. And so we're right in the middle of those two forces. One has a surge of advancement and innovation and then the next. And so what that means at the intersection is it's a little bit analogous to the advent of say electric vehicles where there was a set of infrastructure related to gas motors that is suddenly no longer as relevant.

[09:55](#)

And so there needs to be new infrastructure built to even service this new kind of technology, and that is what we're seeing with AI. It just so happens that if we look at not tacking AI onto the existing financial rails, but thinking of an AI-native way to exchange money, stablecoins are not a silver bullet, but they're very relevant to how machines can effectively and efficiently exchange money. And so there's two forces. While they're both separately getting a lot of attention these days, they're actually quite synergistic and one enables interesting businesses and products and so forth in the other.

Jo Ann Barefoot:

[10:36](#)

We're already seeing people beginning to transact activities with using agents. The financial piece of that is highly regulated, it's highly risky. So what needs to happen in your mind? You speak of infrastructure and we do have the new stablecoin law in the United States, the GENIUS Act, and countries all over the world are enacting legal and regulatory frameworks around stablecoins. But what are the elements of the transition? Who needs to do what to enable the environment where all this is going to work?

Sean Neville:

[11:25](#)

Yeah, so agents, an AI agent today can't get fingerprinted. They can't apply for a money transmission license. And there is existing regulations on our financial infrastructure. And even if AI actors are becoming economic participants, they need to adhere to the existing laws and regulations of all the jurisdictions in which they transact. And so one of the things that needs to happen is we need proper risk infrastructure and mechanisms for AI actors to identify themselves and to say, effectively, "I'm an AI agent acting on behalf of Etsy," or whatever it may be. And you can prove that. And I need the ability to say, "I will interact with that agent that can prove its owned and operated by Etsy, that has these accountability guard rails or these liability protections, but I don't want to interact

with anyone who's pretending to be Etsy or coming from another source."

[12:26](#)

And so this gets to the identity mechanisms. How do we identify these AI actors? How do we hold them responsible if they escape their guardrails when exceptions or errors occur? How can we be sure there's a human in the loop or necessary to handle escalations? And so these are things that don't exist yet broadly speaking. What exists today in most risk infrastructure is an assumption that all bots, that sort of "bots" are bad because they haven't been KYC-ed or there's no legal entity on the other side. And what we ultimately need to happen is we need infrastructure that assumes really the complete opposite. Let's assume that bots or AI actors are ultimately the only actors that will be executing transactions online and we'll be using those AI actors, but we won't be executing the mechanics ourselves. And then so how do we keep out the bad bots?

[13:19](#)

We want to prevent the rogue bad actors, but how do we make sure that we can safely interact with the AI actors that we want to interact with that help us run our businesses and grow our assets and so on and so forth. And so a lot of what needs to happen is even before we get to some of the interesting applications, a lot of what needs to happen is a reckoning at the foundation layer and an understanding of how do we identify these AI actors, how do we hold them responsible? How do we map them back to a legal entity that is permitted to perform certain actions in a jurisdiction and so on and so forth. And then how do we have audit trails to understand why it actually happened and what these AI actors have done so that we can go back and examine it after the fact?

Jo Ann Barefoot:

[14:07](#)

Yeah. This year, you put out an Agent Commerce Kit. Can you tell us about that? Help us understand how this works.

Sean Neville:

[14:19](#)

Yeah, so we found as we were building agentic financial workflows that a couple pieces are missing. We talked about one, which is this identity piece. The other piece is just how do AI actors request a payment or execute a payment once they have an invoice or something or they want to make a purchase? And we believe that both of those things should be solved not by a proprietary solution owned by one particular vendor, but should just work like the internet works, which is we have standards for how anyone can inter-operate. And if you and I create a website, someone else will be able to access that website using HTTP, just using web standards. They don't have to ask our permission. They can just access our content. And so

similarly, we feel that as AI actors proliferate, the identity and basic payment mechanisms ought to be open standards.

[15:12](#)

And so we've contributed the ACK, Agent Commerce Kit as a set of patterns and reference implementations for how this identity and payment flow we believe should work. The idea is that this is open source, it was a contribution to the ecosystem in the community, but the idea is to continue to iterate and evolve these open standards so that we can ultimately build on them. We simply need them to exist. So we can't build Amazon.com if there's e-commerce, if there's no secure way to access a shop storefront on the web. Similarly, we can't have AI actors that execute things like loans or facilitate payments if we don't have these fundamental building blocks underneath them. And so that's what we've contributed, and we'll continue to iterate on it. It's been very well received in the developer community, and so we actually have another release that makes it even easier to use and we'll continue to push it going forward.

Jo Ann Barefoot:

[16:19](#)

I was going to ask that. So you're getting a lot of engagement uptake on it?

Sean Neville:

[16:24](#)

The developer response has been positive, particularly on the identity side. I think those who are working in AI are familiar with the problem, and it's not been difficult to explain. This is a problem that we face. How do we verify the identity of these AI actors? How do we set guardrails and policies around what they're allowed to do, and how do we do a look back afterward to audit them? It's really those three things. And so the reception has been very strong. If this particular implementation needs to evolve, then it will grow and it will evolve. We don't pretend that we've come up with the solution right out of the gate, but it is definitely the right direction.

Jo Ann Barefoot:

[17:13](#)

So you have used the phrase that Catena Labs is or is becoming the first fully regulated financial institution. You use the word bank. Are you envisioning that you will have a bank charter either one way or another, and who will be the customers of this bank?

Sean Neville:

[17:45](#)

So there are two sides to the financial institution that we're building. One side is just pure tech services. So when it comes to things like how does identity work between agents and how do payments happen between AI actors, that's a tech service. We want to make some certain technical fundamentals available to everyone, but as we build on top of those ourselves, we will be executing regulated financial activity. And so the other piece of what we're building is a licensed financial institution, not simply

a tech provider. And the reason that we refer to as AI-native is that we've seen firsthand that sort of tacking AI onto the existing financial rails is problematic for many reasons. One though is that it's quite air-prone, and it's sort of like if you want to build a rocket, it's difficult to take a 1940s aircraft and it's sort of turn it into rocket. Really need to, some of the same flight principles apply, but you need a different chassis.

[18:49](#)

And so we're building things like our risk and compliance infrastructure using AI actors at the outset. As we write software, we're using AI actors to help craft the software. Now, they're not replacing humans, they're increasing the productivity of humans. But when we think about flushing out a compliance organization, people gathering evidence or ultimately filing SARs or whatever it may be, those are semi-autonomous actors in our case. And so we're building agent-native from the ground up in many domains, but ultimately we're also a licensed financial institution. And so the key part for us is if our AI actors are executing things like KYB processes and they're making treasury management decisions, for instance, how do they prove that they're operating under a particular licensure? And whether it's stablecoin flows or traditional fiat flows, how are they executing those flows in a compliant manner? And who is ultimately responsible?

[19:58](#)

Who's the neck to choke ultimately? And if enforcement actions happen, who does law enforcement reach out to get information? Needs to be a licensed financial institution, so we are pursuing licensure. We haven't spoken specifically about bank charters. Conceptually, we are building a new kind of global bank, but we do believe that it is very important to be regulated. Same thing at Circle. As you know, when we were building out our early crypto rails, we went down the money transmitter path and there were many states who felt like we didn't need to be licensed at all. This was so new. We spent a lot of time insisting we did need to be licensed and doing a lot of education for regulators at the time. And now flash forward these many years later and regulators are very savvy about that space. Similarly, with AI, it's early days and there's some explaining, hey, this AI actor is operating on behalf of a licensed FI, but flash forward many years, I think the perspective will be, well, of course it is.

Jo Ann Barefoot:

[21:07](#)

And can you help us picture how these questions will be answered? You said in the agentic situation, how can you prove that you did it right? What will it take to do it right? And will the regulator be using their own AI, I hope, frankly, to evaluate

systems like this? Because I'm not sure how you can do it if you can't use AI tools, but what are your thoughts on that?

Sean Neville:

[21:42](#)

Yeah, so we use something called verifiable credentials now, which is really, it ultimately is a way of leveraging cryptography to provide attestations to others who can verify those things without needing to ask a central source to handle the verification. Simple example would be if I want my university to prove that I graduated, my university can issue me an attestation. And so the equivalent also happens with AI actors. The underlying mechanisms involve cryptography and ultimately signing messages with private keys. Those kinds of things can be fairly complicated for individuals like us to manage or for businesses to handle. But AI actors are quite capable of managing those private keys very effectively. And so the mechanism for proving certain things related to identity is actually very similar to executing payment transactions and stablecoins where we're signing messages using cryptography as well. And so AI can do these things quite reliably, managing the underlying infrastructure. In terms of the examining what happened and sort of looking back, one of the things that large language models are just quite good at is understanding semi-structured or dynamic information sets and parsing them very effectively.

[23:06](#)

And so it used to be the case that in order to accommodate a sophisticated payment flow, we would need a lot of, if this happens, then do this, and if this happens, then do this. And then a record of where the ifs were triggered. LLMs don't necessarily need all of that specified in advance. They can figure it out on the fly. And so when we look back to see what activities and AI executed, whether it's a payment transaction or some other kind of action, we do need large language models to similarly parse this dynamic flow that wasn't necessarily following a predefined set of steps. So there's a little bit of a in-the-weeds answer there, but to your point, if we have large language model workflows, AI actors executing these kinds of transactions and interactions, we also need the same when it comes to policing it and monitoring it effectively, initially as a tool for humans, but ultimately in a semi-autonomous fashion because it'll be much faster and much more efficient and effective that way.

Jo Ann Barefoot:

[24:12](#)

Yeah. As you know, the regulators worry a great deal about explainability as a topic. We see it, it's especially acute in the credit area. It was a legal mandate to be able to explain why you made a decision on a loan, but I don't know if you agree with me, but I feel that there's a big hurdle there in the minds of the

regulators and the traditional risk management community on whether we're going to be able to explain what the AIs did to our satisfaction.

Sean Neville:

[24:51](#)

Yeah, I think that the biggest issue with AI today is it's not necessarily price or even data privacy, sensitivity. It's really, how can you trust these things? And part of building trust is making sure that you can do look-backs and explain why certain activities were done and why other activities were not. And this is even before we get to things related to accountability or liability. If a payment goes off the rails or there's an error, who's responsible for that? Which AI actor is responsible for those things? Ultimately, it gets to the trust issue. And I think there are a couple of things that we have to sort of hold in our heads at the same time. One is anecdotally as we're using our favorite, ChatGPT or Claude or whatever is your favorite, one of these things, it hallucinates a decent amount. It can't give me a solid recipe for a chocolate chip cookie sometimes.

[26:00](#)

And so why would you trust this with your money? And then at the same time, the second thing to hold in our heads is this truly is changing everything, not all at once, but it will change every domain, and all the biggest companies and brightest minds in the world are all in on this. And so both those things are true at the same time. This is as bad as it will ever be. It's going to be very powerful for us. But the big hurdle, the biggest hurdle of all is really related to trust and accountability.

Jo Ann Barefoot:

[26:37](#)

That was actually going to be my next question, because this is fundamentally about whether people will be able to trust this system, the public and also the regulators and policymakers. And to get to that trust, what does the pathway look like in your mind where the regulators are working on these issues, Congress is working on these issues, but we talked earlier about the fact that we have a new stablecoin law in the United States. We don't have much AI law in the United States yet. We've got bits and pieces, but I haven't seen, and tell me if you disagree, a vision coming forth that, we've seen some directional statements of vision from this administration and the last one, which were very different, but I haven't seen the building of the HITRUST infrastructure at work yet. Am I missing that?

Sean Neville:

[27:52](#)

No, I think it ultimately will come. There have been a lot of discussions around human alignment and there may be some regulatory take on that. Right now though, most of the focus has been on more of the foundational layer, so development of models and the hardware that is necessary in order to create these large language models at scale. And so the focus has been

more on that as opposed to say the application layer, and some of the agentic workflows obviously leverage many underlying models orchestrated together. But the trust at that layer is really around solidifying certain processes which now are best practices but are not written to law of how we achieve human alignment so that the outcomes that we get from these AI workflows are aligned with our expectations, whether those expectations are along the lines of ethics or bias prevention or certain specific tools outcomes, accuracy.

[29:03](#)

And there are processes to say if we have an agentic workflow that detects tumors and MRI scans, we have processes to see that the outcomes are aligned with what an actual human expert would expect if they also looked at that MRI scan. And so putting some sort of codified version of achieving human to machine alignment beyond just subject matter expert alignment ultimately will be necessary. And we may see something like that, but I think it at first starts through a series of best practices and frameworks and then ultimately elevates to that level. When will we get to the level of an AI actor actually being able to hold a money transmitter license or a bank charter?

[29:52](#)

That sounds like science fiction today, but ultimately this idea of in all seriousness certification of certain kinds of AIs that are allowed to do certain things if they clear certain barriers, human subject matter expert alignment, testing by an authorized service that executes eval testing. If there's some sort of set of steps that you could do to achieve certification such that we know this AI is certified for these particular activities, then that's a step in that direction as well. It's a greenfield though. Much to be done there.

Jo Ann Barefoot:

[30:30](#)

Yeah, for sure. The two fields, sectors that arguably are the most regulated are finance and medical care and medicine. And I think it's because there too where so much is at stake for people. And secondly, we know that people aren't going to understand that they can't evaluate for themselves whether something's being done right. And so there is a medical, excuse me, is a regulatory framework set up so that the individual trusts that this is a system that is being properly monitored and if something goes wrong, they'll have a way of getting it fixed. So that puts the government in this role of needing to be part of the trust mechanism. The industry has to make the government have confidence that they can trust what's being done, and the government is then helping to provide that trust to the public. Do you agree with that? Even in a decentralized system of the kind you're describing, that the individual will still be relying on regulatory mechanisms?

Sean Neville: [31:51](#)

I do. I mean, I think it varies by use case, but certainly even anecdotally, if we use ChatGPT or whatever it may be to tell us a story, well, the expectation is not the same as if we're asking for how much Tylenol should I give my child? He has 102 fever. The latter is much more serious and it needs to be accurate if I'm going to trust it. And so I think there's a role that is use case or domain-specific. Finance and healthcare are just different than say, generating marketing content or sports and gaming or whatever it may be. And so not all AI actors are the same, just as a lot of the applications that we access are deployed in a cloud system, but a bank that's running on a cloud is just different than another kind of application that's running in the cloud.

[32:50](#)

So similarly with AI agents, and they'll be regulated differently, in the processes to make them trustworthy, we'll have some things that are in common in terms of how we improve these orchestrations. But the specific mechanisms to make them better will be quite different. How an AI actor thinks about structuring short-term, mid-term, long-term memory in one use case is just quite different from another use case. And so ultimately, we're figuring this out now. It's changing very, very quickly, but it's very early days. We're figuring this out now, but ultimately where it'll evolve is there'll be best practices and frameworks and then there'll be certain required certifications for an AI to be deployed in something like a healthcare use case. And then ultimately, that may be codified into broader law.

Jo Ann Barefoot: [33:44](#)

We have a great many regulators who listen to this show because we always do talk about the regulatory side of any topic. Do you have suggestions for them? Some of them might be listening fully aligned with what you're describing, and some of them may be having a lot of trouble even picturing it, not to mention how to regulate it, which some may be thinking about or doing.

Sean Neville: [34:09](#)

So one thing I would say is, I was discussing this with someone else and I thought this categorization of mindsets was helpful. And as we think about the new kind of risk vectors that emerge when AI actors are economic participants, I mean, these are new risk vectors that banking infrastructure is just completely unprepared to handle, and a lot of them are unknown. But as we're beginning to wrap our heads around these new risk vectors and managing, ensuring compliance in these workflows, I think the mindset that is most helpful is the bloomer mindset. So the categories are sort of doomers, AGI will kill us all. The Terminator is coming, this is going to be deadly. And then there's sort of the gloomer mindset, which is actually this stuff isn't that good at anything yet. Where's the ROI? AI is

over-hyped, and all of that. And then there's the Zoomers that are just like, "This is great, let's full speed ahead. We don't need to be so concerned about these frameworks or ethics issues or whatever it may be."

[35:18](#)

And I didn't come up with these categories. I think this may have come from Reid Hoffman, but the category that I think is most helpful in coming up with solutions is more the bloomers. It's early days. We're growing things carefully. Some of those things will not work and they need to be cut and curated, but we're growing these things, and ultimately, the promise is it'll be quite prosperous for all of us. So let's carefully grow these things and keep an open mind and not go too far into the other three categories, and that's the best place to innovate.

Jo Ann Barefoot:

[35:54](#)

That's super helpful. I agree. You said earlier that it's not possible to build a rocket ship on a airplane and that you need a different chassis. What do you think the impact of all this will be on our current financial industry, including the banks? Are these going to reach the point where after a period of figuring out how to do it, it'll become in the tool set and operations of the incumbent industry, or is this going to be too much of a challenge?

Sean Neville:

[36:38](#)

That's a very good question. So going back to the stablecoin side of things, I think it had been the biggest barrier for things like stablecoin adoptions and payment flows in say, take a remittance use case, which is often cited as prime example, where sort of very, a cross-border, fast settlement, low-fee money transmission would be very helpful. The biggest barrier to achieving stablecoin usage in those examples has mostly been lack of regulatory clarity, specifically from the United States. And that barrier has been one that prevented institutions, banks, and other institutions, businesses from doing a proper risk-reward analysis, because it looked like a lot of risk. It was just unclear what exactly a stablecoin was, what are its requirements when it comes to underlying reserves. Does the US government view it as a payment instrument or is it some sort of money market fund?

[37:40](#)

It was very unclear. And now that's clear. It's at least clearer with regard to the US dollar stablecoins. And when we say stablecoins, we mostly do mean US dollar stables at this point, and the desire is to move dollars. There are other stablecoins in other currencies, but mostly when people refer to stablecoins, they mean being able to access US dollars globally. And so now there's that, there is that clarity. And so the next hurdle becomes an infrastructure one. How do governments, banks,

other institutions incorporate this new infrastructure that is required to actually execute stable coin flows? And how do they plug it into their ERP systems, and how do they use their existing treasury management or supply chain software with these dollars, but dollars that are custodied quite differently? And similarly with AI, there's an infrastructure challenge there. And I think some of the approaches have been, well, let's keep the core, let's keep the chassis, but we'll add it onto the edges.

[38:43](#)

And when it comes to AI, that becomes problematic because it's very difficult to make AI execution of the existing tools reliable. And because we have these machines that are non-deterministic, it can handle very sophisticated inputs and make decisions in some cases, but it doesn't always do what we expect it to do. And so when we have these existing systems that are executed via a tool or an API and work a certain way that is not designed for AI, then it becomes more error-prone and problematic. And so we have an infrastructure challenge there as well. I think some institutions will be able to tackle it through partnerships and as well as growing their own technology. Others I think will struggle, which is of course always the case.

[39:33](#)

Ultimately what we came to with Catena is the belief that in order to have AI really effectively move money and safely custody it, we needed to start from a foundational level and just rethink a lot of this, the existing software, and not necessarily assume that we'll be using anything like credit card rails if they're AI actors. They don't need to have credit card systems that pass through this number, a three, five, nine, 12-party system in order to move money in the same way that e-commerce flows between businesses that are designed for humans. And so we're tackling the infrastructure problem just from the foundation, we're going to build up from the foundation. But others will be able to incorporate the infrastructure in different ways. But I think ultimately with regulatory clarity at hand, the big challenge becomes an infrastructure one.

Jo Ann Barefoot:

[40:37](#)

When you think about this future, one of the benefits to the financial consumer will be lower costs across the board, which should make services more affordable. Do you agree with that?

Sean Neville:

[40:55](#)

Yes.

Jo Ann Barefoot:

[40:56](#)

More efficiency.

Sean Neville:

[40:57](#)

That's right.

Jo Ann Barefoot:	40:58	And what are the other benefits to the end user of financial services, including on the retail side?
Sean Neville:	41:07	Yes. So the real value of stablecoins by themselves are really, they're two fundamental things. One is the speed of settlement. It takes milliseconds to settle in some cases. And then the other is cost. It's almost free. It's not free. It's almost free. It's so inexpensive that in most usages of it, it can become free, similarly to sending an email around the world. Emails don't stop at borders, and money shouldn't stop at borders. It does actually cost a little bit of money to send an email, but it's so fractional and the cost can be absorbed in other business models by people who manage sending email for us. Google has an ad-supported model, so they give Gmail away for free, that sort of thing.
	41:49	But those are the two fundamental things, speed of settlement globally, anywhere around the world, and very, very low cost. On the AI side, we have similar advantages and we'll ultimately get to the point where AI, semi-autonomous working for our businesses and for us as individuals will be the safest way to manage our money and also to get a return on our money. And it can seem strange now in our experiences with some chatbots, but when we're talking about agents beyond the chatbot era, ultimately they will get to the point where you just wouldn't trust any other system other than sort of a certified, regulated AI workflow to manage your assets and to get a competitive return on the assets that you're holding.
Jo Ann Barefoot:	42:36	Yeah, they're going to be able to know things that we don't have the ability to know. We don't have the time to find out or we don't have the information available to us in accessible ways, which over time is going to make the system, there's so many problems in finance that are about opacity and complexity and people making judgments without full information. And in general, the AIs are going to really help us see into these systems, I think. What about financial crime? The doomers are very concerned that crime will worsen, and we certainly see AI crime in finance on the rise. We at AIR just did an exercise that we called the AI battle front on fraud, and yet the same tools are going to equip the system with better ways to root out such things to detect them. Are you an optimist on that?
Sean Neville:	43:45	I'm ultimately an optimist, and the way I think about it is not necessarily AIs replacing humans and magically doing everything better, but starting as a tool that becomes a much more effective tool for humans to leverage as they're monitoring for financial crime. And so I'm very optimistic about it. I think in

terms of transparency and understanding why an AI workflow made certain decisions, there's a lot there. And the way that a lot of our systems fit together today are based on APIs or certain integrations that are very much intended for humans to be able to understand what's happening. And so if you think about AIs talking to other AIs, will they necessarily use these things? Do they need an API? If they're talking to one another, would they even really be talking in English or any language that we understand except as a way to provide some transparency into what decisions are being made?

[44:54](#)

And so I think there is a lot there in terms of truly optimizing AI to AI communications. And I should say, by the way, when it comes to payments, we don't see a lot of AI to AI payments yet. We believe it's coming, but it's more like AI systems paying for access to resources or accessing certain services and paying for those, or paying humans in some cases, or humans paying into AI workflows. Not a lot of AI to AI yet, but as that begins to emerge and there's more AI to AI communication, how do we make sure that it's transparent enough that we can understand as humans why decisions are being made so that it is an effective tool and that we'd be able to use it as such for policing financial crimes, but also for other sensitive activities as well. That's really a fascinating space.

Jo Ann Barefoot:

[45:42](#)

That's for sure. Do you have any other advice for players in this space, people who would be listening to this podcast, whether it's incumbents or policymakers or consumer advocates or law enforcement, people looking at the financial crime, money laundering, fraud scam challenges? Anything that people should take away and try to think about differently or do differently from where you think?

Sean Neville:

[46:13](#)

Yeah, I think in addition to the mindset sort of comments about kind curiosity and controlled optimism about this, I think more specifically, it's being able to get comfortable experimenting with and exploring agentic tools. This is a step beyond using chatbots, but using and getting familiar with AI tools and some of the issues that come with using those related to evaluating them, how guardrails are sort of put in place and defined. And that side of AI, it's not as commonly explored as say the latest model to generate a video or something, but really around how do you apply policies to agentic workflows and what's being done there now. And there are major vendors that are working on this. Anthropic is a leader in this space, OpenAI and others. But in terms of digging into material, I think that's the space to explore. So ultimately, understanding these guardrails and how they function and where they fall flat is key to being able to

police AI activity effectively and understanding where the risks are emerging and ultimately also how trust can be built.

Jo Ann Barefoot: [47:43](#)

I can't thank you enough and I really enjoyed this conversation. Is there anything we haven't talked about that you want to add?

Sean Neville: [47:50](#)

I think we covered a lot of ground. One thing I would say is that six, seven, eight months ago between those two forces I mentioned, so the interest in AI orchestration, almost hesitate to use the word agent because the joke is that no one can define what an agent actually is. But these AI orchestrations or workflows and then the other force being stablecoins, really the surge in innovation had been very much an AI orchestration. And there is a lot of innovation happening there, but not so much on the stablecoin side of things. In fact, as some of the early cases that we worked with, the financial use cases that we worked with in AI, to the extent that we were leveraging stablecoins, we didn't necessarily highlight it. We referred to it as digital dollars. We made clear when we were using USDC that it was known, but at the time, the environment was quite different and the perception of stablecoins was quite different in that it had not yet achieved escape velocity with regard to the rest of crypto.

[48:59](#)

And crypto in many AI circles still has a fairly toxic reputation. That has just completely changed over the last few months, and the surge of interest in stablecoins is unlike anything we've seen. And I've been involved in crypto for 14 years now and started Circle 12 years ago. This is a new era of interest, and stablecoins, I don't know if we keep calling them stablecoins. Maybe it's just money or dollars or something. But the surge of interest in tapping into dollars globally and leveraging these public blockchains and these public networks to do it is just, it's outpaced anything else over the last couple of months in terms of the interest that we're seeing on the finance side.

[49:48](#)

So both of those things are still strong, but I would say that one of the things that has changed over the last two, three months is this interest in how to leverage the capabilities of digital dollars that we're calling stablecoins. And some of this is really not so much about the technology. It is faith in the dollar and the US dollar and wanting to be able to use the dollar all around the world as inexpensively and as safely as possible. So it's maybe not news, but it is a change and a shift that has happened.

Jo Ann Barefoot: [50:22](#)

Yeah, that's a great insight. We at AIR are doing a lot of work on stablecoins, including starting a project in Africa where there's so much interest in the potential to use them, especially for

cross-border payments, but other use cases as well. Sean Neville, I cannot thank you enough. I have really enjoyed this conversation. Where can people get more information about Catena Labs?

Sean Neville: [50:50](#)

So we're at CatenaLabs.com, and we have also released this Agent Commerce Kit, which is at AgentCommerceKit.com. And so that's freely available. We have been somewhat quiet about this specific license commercial product. That is forthcoming, much more news about that later this year. But meanwhile, would love to hear from people who are working in the space and we'll share much more as we move forward.

Jo Ann Barefoot: [51:20](#)

Wonderful. And we will put more information as well in the show notes at regulationinnovation.org. Thank you for joining me today. It's been wonderful.

Sean Neville: [51:29](#)

Thanks so much. I appreciate it.