

Agentic AI in Banking & Financial Services: Beyond the Hype

VIDEO TRANSCRIPT

Jeff Veis: I'm Jeff Veis, Chief Marketing Officer of Impetus Technologies, and I want to welcome you to a very special experience we have in store for you. One where we're going to focus on real world impact, real world issues as they relate to agentic AI applied specifically to financial services.

We're going to try to go beyond the hype, beyond the obvious information that you're being exposed to on a daily basis, and talk about what is really happening, what financial services companies are doing with regard to Agenda AI.

Both where they're succeeding and where they're delivering measurable results.

Likewise, where they're being challenged and how they're addressing those challenges.

Our goal is to hopefully share with you key insights, actionable insights about what you can do within your own enterprise, regardless of size, to harness the full potential of the Agentic AI.

Now, first, I'd like to welcome our distinguished

guests, Stephen Walker from Global Data and Deepak Khosla from Impetus Technologies.

Can I ask you gentlemen to both introduce yourselves, your role and your respective companies and then we'll get kicked off? Why don't we start with you, Stephen, right?

Stephen Walker: Yeah. Thanks. Jeff. Delighted to be here. So, Stephen Walker, I'm a retail banking analyst based in London.

I've been doing this for 15 plus years time at Forrester Gardner Lloyd's bank here in the UK. But the last 8, global data where I look at the impact of disruptive technologies on FS providers.

So to December last year, we acquired Celent. So, we're now very excited to combine that deep sector specific tech expertise from their analysts with what we've long done within GD FS using consumer SAT data, executive surveys, forecast market shares and so on to give our clients, our technology vendors and retail banks an

extremely multi-dimensional sense of their performance and how they can close and grow gaps against their competitors.

And obviously these days, everyone wants that extremely multi-dimensional view focused on AI.

Jeff Veis: Deepak, could you introduce yourself?

Deepak Khosla: Absolutely delighted to be here. I'm Deepak Khosla joining in from Bay Area, California. I lead the growth office and the AI business at the impetus.

So, at impetus, we work with large enterprises, you know absolutely with the financial institutions on data modernizations, AI ready data foundations, context engineering, agentic AI solutions and so on so forth.

Our perspective actually comes from the implementation side, you know, from the expertise Assemi side, where AI has to move beyond the demos and the POVS to basically operate in real enterprise systems, which demands data controls, lineage, regulatory expectations and so on so forth.

So, that's what we at impetus do and happy to be here and talk more about how AIS shaping our financial services world.

Jeff Veis: Very good. Well, let's, let's kick it off and dive right on. You know, we see in adoption surveys today, they're citing numbers that essentially anybody and everybody is looking at experimenting with and deploying Agentic AI. So, that that's the norm now.

But we also hear through the grapevine, through formal surveys that only a relatively small fraction are actually in full production. And it's they certainly get lots of visibility, lots of buzz around the industry. But that gap has been stable for a while now.

So, I want to ask you, well, this dichotomy between all the hype, all the support for it, but this relatively smaller set that have moved in production, is this a temporary bottleneck? And if so, is it because of data readiness, is it governance, is it model risk policy? Or is it evidence that there's something unique in financial services that makes it fundamentally harder to use agentic AI to apply it to compelling use cases that deliver just what's your viewpoint?

What are you hearing? What are you observing? And Deepak, why don't you kick it off and then we'll bring Stephen into this question as well.

Deepak Khosla: Absolutely. See, I think it's both temporal and structural. Temporal because teams are still learning, organizations are still learning how to build and govern these agents. But structural more importantly and specially in financial services industry because this industry is very different from other sectors, right? And agent in banking is not just only summarizing a document, right?

It could actually influence credit, it could influence a decision on a fraud payments, you know, how a customer is treated, regulatory reporting or for that matter, financial advice, right?

So, the bets on production or the bar on production is much higher when it comes to financial industries, banking industries versus any other industries, right?

So, both things need to be kind of, you know, looked at together.

Now, I mean, at least, you know, from our viewpoint, the customers that view, we basically

work with a lot of these customers and Jeff, Stephen have gone into production with some use cases, right?

And we'll talk about it in a minute. These companies have kind of, you know, tackled few key blockers, which I think they have done quite well. At the same time, you know, these are the blockers which, you know, some of the other financial institutions are struggling with. What are those blockers? In my mind, there are these five, right?

1. The data readiness, you know, we all know in, in banking industry and financial services, you know, data is fragmented, you know, across payment risk, CRM, case management, legacy systems and so on so forth. So bringing, bringing a unified view is one issue.
2. Second is the context cap. You know, and if you've been falling input is we've been talking about this a lot, but you know, what we feel is the agent needs the business meaning, not just the data access. So, they need the meaning behind the customer, the product, the policy, the risk, the compliance, and so on so forth. So, the context cap is the second one.
3. Third is governance, right? You know, banks would need audit trails, they would need approval gates, they need very well defined entitlements, they need defined and provable model risk controls and what kind of human accountability is coming into play. So, governance is the third major blocker.
4. The fourth one is integration, right? So, you know, agents do not work on their own. They need to connect to the real business workflows, answer questions, and so on so forth. So, how these agents are integrated across financial institutions, banking systems is key.
5. And the last and the most important thing which is I would say the final barrier is the liability, right?

So banks are still kind of defining, you know, what an agent can recommend, can assist, can escalate, can act, and what are the liabilities behind it.

So these are in my mind, you know those five key things, which the institutions that we are working with have taken care when they move to production and the other institutions we are still grappling with. Stephen, your viewpoint?

Stephen Walker: Yeah, thanks, Deepak. I would agree with all of that. I want to add that I think it's a pretty confusing time in the marketplace to work out actually what is going on, right? So there's constant reference to agents without really defining exactly what is meant by that. Often some tech vendors would talk about identity capabilities that we might think are more accurately described as basic bots, right? So their element of agent washing from the other end of things, you've got complex multi-agent orchestration, which is of course much harder, and it takes much longer to do.

So, that's always going to take longer to move into production. So, we haven't got these apples to apples comparisons. I think you've got to start with a level set on the definition. So, to be clear on the precise unit of analysis.

And then you can look at how common it actually is, not just in ambition, because it's everyone's ambition, but in actuality in production.

So, if we take a simple definition of an agent as they can process information, make a decision and implement something, say there's not an enormous amount on that - point one. And this notion of pilot graveyards or pilotosis is well documented.

I think in terms of what causes that right there are many things. From the banks that we speak to that are having the most success moving from pilot to production and having diminished marginal cost from each new use case that they deploy. They've managed to create a unified data access across all channels functions in near real time.

The first thing that happens when you're trying to do a new use cases is where can we find the data? That's the thing that always takes the most time. Banks that that address that head on. They create operational advantages that compound over time.

Of course, one of the challenges around tackling that legacy though, it's very institution specific, right, about what to address first, when and how depending on a on a variety of factors. But yeah, that would be my top level response there.

Jeff Veis: So interesting. So, you know, there are multiple variables that all organizations, but especially banks have to deal with because the bar is higher. Well, one you brought up was compliance and policy. So, let's double click on that.

You know, from a regulatory posture, you know, regulators have been signaling caution, right? That's kind of what they do, right?

They said rather than writing hard rules here, you know, updating model risk guidance rather than creating agentic AI specific framework. So, this whole idea of do you capture it in the framework, or do you capture it at kind of a higher level?

And the question to you guys and maybe I'll start with you, Stephen, is this ambiguity that maybe

that's out there now of where do you address this?

Do you hard code it?

Do you soft code it?

Do you have a combination?

Does that favor, and I'll give you maybe a false premise, but does that favor the large incumbents with massive IT organizations, massive regulatory groups that can over engineer compliance because they've been doing it so long?

Or sometimes when we see a new technology and we see an inflection point and a disruption, it creates an opening for smaller, faster mover players.

We've seen this occur in multiple industries going back to when .com came.

So, if you have the small guys and the big guys and everybody now has access to the same models, right? Anybody can get a subscription to Anthropic, for example, is there an opportunity or an advantage of one over the other regarding kind of working through this? How do you apply regulatory policy and enforce it when you're injecting Agentic AI into the fabric of a bank?

Question to you first, Stephen.

Stephen Walker: Yeah, thanks, Jeff. Maybe start with a sort of high-level observation around some of the business challenges there. So, clearly the regulatory landscape is tricky because it's evolving. There's uncertainty, these inconsistent expectations.

And, and as you say, in many cases, there's not

detailed AI rules. It's left up to banks to interpret how to imply those rules. And that lack of guidance has meant developing their own ad hoc frameworks and governance processes and so on, which big banks have an enormous amount of experience with.

And they have that to protect consumers and to protect the institution. Now I'm struck by let's take an example from America, Bank of America. We're probably one of the best bots on the planet for an incumbent bank with a 98% containment rate with Erica.

But their tech leadership have been adamant that there's no use of generative AI or agentic AI without a human in the loop in the near term.

Now that creates space for new entrants that might be able to move more quickly.

And what's the evidence that that we've that we've seen on that? I think examples in Europe.

So, Bunk, a digital bank in the Netherlands was early to get sign off from the regulator to use artificial intelligence for fraud detection. They managed to prove that it was 20% more effective than prevailing methods. And then more recently, their personal financial assistant has got lots of use of agentic contingent to AI thin, which has had very positive customer feedback.

Yet at the same time, they still had fines from the regulator, but overuse of AI for onboarding and AML and some KYC violations.

Another example is Starling Bank in the UK, leading digital bank, excellent customer sat scores and they brought out a new personal financial assistant again using agentic capabilities,

but they've also been fined by the regulator in the onboarding context around AML and KYC.

So, I think there are real challenges for younger digital banks that might be able to move quickly with a marketing win of being first that may cost them fines and may backfire around this to trust with this technology more medium term.

Deepak, let me let me pass that to you.

Deepak Khosla: No, no, I agree with you. I absolutely agree with you. I don't think so regulatory ambiguity, you know, kind of favors either a large bank or a smaller player, right?

I feel it would favor institutions who can turn ambiguity in a little bit more of a disciplined operating model by leveraging the agentic A capabilities that we have now, right?

Of course, large banks have a little bit of an advantage because they already have some mature model risk compliance. You know, all those control functions kind of are well defined and listed and working for many years, right?

They can also absorb, you know, the cost of over engineering if they want to do as well as the governance.

But at the same time, I think you know, to some level agentic AI and the advancements in the models and you know, what the labs are producing has kind of level field the play bit more, right? So, I feel advantage will go to the firms who would do two or three things quite well.

1. Have a very clear internal policy of the agents before even regulators actually force it. That how the agents will actually work, which

essentially means kind of define what agents can access, what they can recommend, what they can prepare as responses to those regulatory needs approve and execute, right? And that's number one.

2. The second, you know, have a little bit of controls in the architecture. These AI models would not know how your institution actually works. So you know, what kind of lineage do you need? You know, what kind of water trail do you need? The human approvals, you know, evaluation, monitoring and so on so forth. Think about it at this at the start itself, not kind of bolt it on later.

If I think if I think these kind of things is done well. So, the if the guardrails are built well, then I think the rewards would be reaped by the prepared institutions.

It's not about a small or big one.

I think we are clearly seeing and now we are seeing across not only financial institutions, we are seeing in a lot of other businesses how organizations with only two or four or five employees can turn a billion dollar revenue in no time, right?

So, from that perspective, this is this kind of, you know, game changing.

Jeff Veis: Interesting, interesting. So, it sounds like there's a bake off, right? It's not clear if there's a winner, if there's multiple winners, but there's advantages and disadvantages to these, these being small and nimble versus being large and resource capable.

Well, let's take that one step further.

Let's talk about what I'd like to position as this funky, to use a technical term, trust paradox that you have.

We've all gone to our favorite LLM, and I've certainly witnessed it where the LLM wants you to feel good, wants you to persuade you that what you're thinking or what it's telling you is the right way to go. This kind of conversational engagement.

So, rather than giving you this hard right answer, and I guess if you set this up as a trust paradox, and I'm going to start with you, Deepak, there's an argument that these conversational AI agents are more natural way for humans to build trust with these rigid banking financial services interfaces.

That's part of the attraction - is I can feel like I'm having a conversation even though I'm getting very serious information about my financial condition or investment decisions, etcetera.

The question is, or the premise is this conversational fluency that makes an agent fundamentally more persuasive.

What if it's wrong?

What if it's telling you sell all your XYZ shares when that's the wrong way to go?

Or you're saying, should I sell my shares?

And it's saying absolutely, because I want you to feel good. I want to give you that feel good.

And you know, is that possibly a risk?

Is that more exploitable from a kind of a social engineering hack or does in theory the

intelligence and all the forces behind these models, you know, give you the right answer?

Deepak Khosla: No, I think the trust paradox is real, definitely it's quite real. You know, we all know that agents can answer any kind of question. It's not that they cannot. They can answer any kind of question. Conversational AI can actually, you know, make banking applications while more accessible and personalized.

But sometimes, you know, that this kind of a fluency can create some false confidence, right? Customers may trust, you know, any answers because it might sound quite authoritative. It might sound, okay, it's got some more reasons behind it, or it might sound that, yeah, this is what I wanted to hear. This is what I wanted to do.

You know, the agent is actually telling me exactly the same thing.

So, it might give, you know, kind of a false trust to the application, right? What I think key over here is the underlying context, right?

The key over here is having, providing the context to the agent, which, you know, comes from the various applications, various data sources. It comes from the continuous learning that some of these agents would have had.

The continuous memory that some of these agents would have, would come from a little bit more of personalization, which is coming in based on what customer is kind of, you know, I'm interacting with a little bit more of background with that.

I think if lot of these things is taken care of, then I feel we would have the right set of grounding,

right set of explainability, and a light set of controls in place for these agents to be a little bit more trustworthy and for customers to feel good, right?

The other key thing that you talked about was human in the loop.

I personally believe and that's what we are, we are also talking to our customers. Human in the loop is not only an intervention, it's a model that needs to be think about right. You would need to think about various kind of workflows.

You know those workflows which is more often a low risk information providing, to a little bit more of a decision making to a little bit more of a regulated action that would be taken care of.

If that happens, what is the human role in it, right?

So, maybe when it comes, you know, to information sharing, it's more about, yeah, just sampling and monitor it.

If it's about some kind of recommendations, then you know, you start looking at exceptions and, and see based on those exceptions, how the agent is actually performing, right?

But when it comes to let's say some kind of a regulated action or a regulatory workflow that needs to be taken care of, there you need to have a mandatory review and audit, right? There you need to have the right set of escalations path in it.

So, I feel these are few things which will make some of these agents perform better and take care of the trust paradox that we had.

Jeff Veis: Let me bounce to just a quick spotlight

on Impetus. Deepak, you're the one to ask this. Some people know Impetus.

We have some major clients in the financial services space, but to some it may be a newer name.

Can you talk about not just who Impetus is, but in terms of agentic AI, which the company is all in on and specifically financial services, what methodology you know, every, the technologies are all out there, right?

So, if that's kind of a level playing field to some degree in terms of access, what is Impetus's unique approach or framework to try to help financial services institutions think about adopting agentic AI, scaling agentic AI, managing agentic AI?

What use cases are we seeing today to give people just a brief idea of who is Impetus and why is financial services so key to us?

Deepak Khosla: Yep. No, absolutely love to answer that question, Jeff. You know, our approach starts with the belief that agentic AI success will depend on quality of enterprise context. And in the financial services world, based on, you know, all the customers that we have worked with, not only in North America, but also in Europe and Latin America, it actually means context.

It means, you know, bringing the customer data, the account data, the regulatory obligations, the approval thresholds, risk policies, all that together, right?

If that particular context is fragmented or it's poorly governed, agents will fail in production.

Now in order to do that, we do couple of things.

Number one, you know, our customers meet us at different point of time. There are customers who are still in a place where they want to basically build this whole AI-ready data foundation.

So, when we do that, when we help them modernize, migrate their data estates to a more AI-ready or a cloud-ready foundation, we make sure that the context carries along with the data. If you just don't move the data, we basically bring the context and we help them select the right set of, you know, data sources. Also, it's not only about structure data, you know, the context actually sits in unstructured data.

So, what is those particular sources and, and scenario looks like? So, we help them basically move, migrate and create that AI-ready data foundation. So that's the customers who are still in the process of building that particular layer.

But then there are customers who either have agents who are not working or who have the data all sorted, but they now want to basically monetize this whole AI wave, right?

When it comes to those customers, we help them build the whole context layer. And for us the context layer means, you know, building the semantics, the ontologies, the knowledge graphs, the rules and the policies and having a very good memory kind of build in, right?

So that's the context layer that we would build in.

it's a core part of our engineer set of offers that we basically give to our customers, right?

Now, once that is done, then comes the hard part

of building those agents, right?

So, you have built, you have moved the context, as you build this whole AI-ready data foundation, then you have built the context layer. Now you have to make sure the agents use that particular context layer, right?

And that's where we basically help looking at different business processes, looking at different sub processes.

Sometimes it's only about, you know, managing the false positive. Sometimes it's about KYC, sometimes it's only about BCBS or you know, some kind of regulatory compliance, various set of use cases.

Then we create orchestration agents, right?

Those are the agents which are kind of built around those workflows. It could be sub-activities, or it could be, you know, larger domain activities, like fraud investigation, dispute resolution, we do personalization for credit card companies and so on, so forth, right? So that's the third piece that we help.

Fourth, what we do is while you have now you know, we have, we have created your data foundation, you have created the context layer, you have, you know created those agents, now you need to govern those agents.

So we have solutions, we have platforms which help evaluate the outputs. Monitor the drifts which is happening both in terms of context and the response, have an ability to basically control the excess, figure out the audit trails and more importantly the token consumption.

So, all that tokenization is something that we take care of. This is how we help customers in various form and factors that come into play.

Now we also believe that there has to be a very different methodology and engineering life cycle to basically deliver what I just said, right?

And over there, what we have institutionalized, and now this has become the religion of Impetus, is CEDL, right?

So, that's an engineering approach that we follow for all the work that I just talked about.

CEDL is short form for Contextualize, Engineer, Deliver, and Learn.

Just make sure that every project that we do, at every point that customer meets us, we make sure that the context is captured, the context is engineered, the context is delivered via EDD methodology and so on and so forth.

And finally, we learn how those agents are performing and we bring the agent, the learning back into the whole context layer and, and make sure those agents perform better, right?

So this is how Jeff, we help our customers in four different ways and with methodology, which is called CEDL.

Jeff Veis: Well, thanks. Well, it sounds like it's end-to-end. And it's not just about deploying the technology, it's not just about getting your data organized. It's really about understanding what's unique.

That's what has struck me is agentic AI is ubiquitous in one sense, but incredibly unique

because your context as a company and your customers is unique. It's a snowflake, if you will.

And while the technology may be common, the approach and how you put it forward is going to be very, very specific to that company, their business objectives, their risk tolerance.

And that melding can only really be done with the talent, either internally, externally, or combination to bring that together, because it's not deploy the software and you know, it turned it on. It's harness that technology but understand it on how your company does its work, which is, you know, is moving target.

Let's get to the last question I want to put to both of you. I'll start with you, Stephen, which is: now what do people do?

What is their next step? What advice would you give to people either in enterprises that are very sophisticated, that might be listening to this session today. That are in the mix with agentic AI, that have the technologies running, have the expertise, have the data scientists. What advice and guidance would you give to them based on your perception, your experience?

And then in the same way, what advice and guidance do you give to people that probably have heard about it, but maybe they're not as far down. Maybe they are in the earlier stage, and they want to not repeat maybe mistakes that first movers have made. As a middle mover or even a laggard, what advice would you give to them?

So, if, you know, when people wake up tomorrow, what should they maybe think about doing to be able to, in a smart, intelligent way, harness agentic AI in the financial services space?

And Stephen, question to you first.

Stephen Walker: Yeah, thanks, Jeff. Let me take the second part of that first. And I'm mindful of what you said at the beginning about saying something beyond the obvious, but I think it's disingenuous to say anything other than what I'm about to say in answer to the question as you as you stated it.

So, for that, for that bank that is less experienced, less mature, I think the number one recommendation and this comes from all the research interviews we do with the conferences that we run with CTOs from various banks, the need to take a calm and measured approach, right?

There's obviously an enormous amount of preparatory activity around this one. One way that we measure that our global data is through our signals data.

So, we look at things like hiring activity, painting applications, M&A activity, corporate filing mentioned tag to specific technologies. And the volume and the velocity of increase that we see for AI outpaces all the other technologies we cover.

I'm sure that will be a surprise to no one listening. But we also see that same trajectory now within AI, which to say that agentic AI is far outpacing generative AI and machine learning algorithms and so on.

And so with that has come this real overkill issue that every new vendor product has got an agent at the front.

Everyone in the bank wants to use an agent for a new task. And but undeniably, right? There's enormous potential here.

I heard CTO describe it as a gift from the gods in a conference a couple of weeks ago.

But that doesn't mean existing approaches no longer matter like traditional automation or the traditional rules and ROI no longer apply.

And you can get to this slightly absurd scenario using agents for everything. A colleague quipped that it's like taking the kids to school in a Formula One car. So, I think they're being measured and calm and thinking from a tactical time for market perspective, what is the next process?

Whether there is genuine consumer pain and sits at a sufficient level of complexity that we can apply an agent to that task and drive some concrete business results.

And then medium to long term, it's all those questions around the architectural changes that Deepak had talked about, and I talked to them briefly in the upfront answer.

Let me pass the more advanced end of that question to Deepak. So, the first part of it.

Deepak Khosla: My advice is that let's not start with asking what agents can do, right?

I think every bank is a decision factory, be it credit, be it fraud, there is reception services, literary needs, KYCS and so on and so forth, right?

So, there are thousands of decisions that would be happening on a minute and an hourly basis.

And agentic AI is largely a starting point or decision technology, which is an inventory of the decisions that the bank will take and not some kind of a tour of, you know, art of possibility, art of science or what AI can do. You know, that's something AI labs can do best, right?

So, in my mind, you know, there are these criteria's or a three simple test that, you know, we basically look at when, when we go to banks and I'm kind of simplifying it. There's much more to it. But number one is volume.

You know, is it something which is, you know, thousands of transactions that happen or the thousands of you know, similar activity that happens every day.

Second, if we find a way to automate it, if we find a way to basically do it better, if we find a way to make it, you know, better for the consumer, is the ROI measurable?

And the third one is, is the whole process documented? Are there written rules? Is the context already ready or not?

This actually helps a lot in terms of identifying the right set of use cases. This helps a lot in terms of, you know, identifying where the focus should be rather than doing as I say, you know, those technical POCs which will always give you good results. But how does the business actually use them?

Jeff Veis: Okay, well, I appreciate that. It sounds like there are ways to be able to intelligently get started and maximize time to value. It certainly doesn't feel like, you know, everything's figured out.

There's probably a lot more learning and discovery and experimentation to occur in the future than has occurred yet.

And this is a tall statement. We've seen technology trends come and go.

Some were over hyped and didn't deliver. Some were really, really transformative. Let me throw one last open question, one last kind of take away.

If you had to impart one thing to kind of put the cherry on top of what we just discussed, what would be that one thing you would impart to the audience today about thinking about agentic AI, about being successful with agentic AI and being able to do it in a fashion that is managed, is thoughtful, is intelligent.

And let me kick it to you, Stephen, one thing that you would put a capstone on this discussion today.

Stephen Walker: Yeah, I think as you alluded to there, executives that are 20 year careers in banking have already managed through a number of different hype cycles of cloud-native, mobile-first, digital-first, .com. There's an element of this movie before and experienced practitioners need to be sort of impervious to the hype and make sensible evidence-based decisions. I think that the number one thing is not to spend the most on AI, right?

That is not the goal.

The thing that is unchanged is solving consumer problems and long term profitability.

So, with that, we start with the consumer pain and then we work into how agents can play a role in

that solution.

And then I think for big incumbent banks, the key thing is when a solution has worked and it's been proven, is rolling out that capability across all of the processes that can leverage that that capability to drive as incremental cost reductions over time.

Jeff Vies: Yeah. Deepak, last microphone to you, what is that one gem of wisdom you would impart to the audience today?

Deepak Khosla: Yeah, I think we talked a lot or rather I talked a lot about context.

You know, I feel context layer is proprietary by definition. It's something that you cannot buy. It's something you have to engineer it, right? And hence, you know, my advice would be to basically, you know, think about how to build that particular context there, which becomes your proprietary mode over the models which almost everybody will have access to, right?

So, you know, think about what does that context need?

What does that agent need?

What kind of data sources it needs?

How can that context be trusted?

How can it be accessible to all the agent and agentic workflows?

You know, we talked about authority, you know, what kind of authority these agents will have depending on that particular context and how do we prove that, you know, the actions that the

agent has been taken based on those contexts is right.

I think all that architecture is something that needs to be built on, which will make, you know, agentic applications proprietary rather than, you know, just working on the world models.

Jeff Veis: Well, I love that. It's everything and it's simple in that all roads lead to context and context is going to be not just an enabler, but what's going to differentiate you because everybody can get that same model, but only you can have your context.

Well, thank you both. This has been a great conversation. I think it's fair to say we've gone beyond the hype and that was our goal today.

And we'll probably need to do this, it feels like every week. But realistically on a regular basis, because this space is moving so quickly, it makes it both interesting, exciting, but also challenging.

So, thanks for the engaging conversation today.

And to the audience, thanks for your time and interest. We'll keep this dialogue going and exciting things ahead for all of us. Thank you very much.