

Episode 149 | Microsoft deploys the FDEs (Kevin Shatzkamer)
Released September 15, 2026

Kevin Shatzkamer: [00:00] I think we're at this point now where we've said the bottleneck is no longer an AI model, but I also think it's not cloud, it's not data, and it's probably not even engineering capacity at this point. The bottleneck is really transformation capacity, and that's Microsoft's focus on what we want to bring to telcos.

Announcer: [00:23] This is "Telco in 20", a podcast that helps telco execs achieve a competitive advantage with AI in the public cloud. It is hosted by Danielle Rios, also known as DR. Today, we're talking to Kevin Shatzkamer, corporate vice president leading the Telco and Media Engineering Studio for the Microsoft Frontier Company.

Danielle Rios: [00:42] Hi guys, I'm DR. Every telco I talk with has AI everywhere, except in production. Hundreds of use cases, dozens of pilots, thousands of agents. Ask why they're stuck, and you'll hear about the usual suspects. The models aren't ready. The data is a mess. We can't hire the talent.

[01:03] Microsoft's take, it's none of those things. The bottleneck is transformation capacity, how fast you can turn an idea into production. Here's why that's hard. AI doesn't break in the demo. It breaks at the seams where a model hits fragmented data, decades of technical debt, and workflows that cut across the whole org chart.

[01:24] So Microsoft has launched a new operating business to close those seams, the Microsoft Frontier Company. It's not a product, it's engineers embedded inside telcos doing the work with them. Today I'm joined by Kevin Shatzkamer, who's leading the Telco and Media Engineering Studio for the new group. We dig into where AI stalls inside telcos, why Frontier runs on two week sprints instead of two year programs, and what telcos bring to the table that Microsoft can't manufacture. So let's take 20. Kevin Shatzkamer is corporate vice president leading the Telco and Media Engineering Studio for the Microsoft Frontier Company. Hi, Kevin. Welcome to "Telco in 20".

Kevin Shatzkamer: [02:09] It's great to be back. I'm excited to have a conversation again today. We got a lot going on.

Danielle Rios: [02:13] Oh my gosh, I'm so excited to talk about helping enterprise companies get their AI up and scalable and running.

So we're going to learn all about this. Microsoft recently launched a new operating unit, the Microsoft Frontier Company, with engineers embedded directly inside enterprise customers to build and run AI at scale. And so why did Microsoft think this is super important?

Kevin Shatzkamer:

[02:36] It is. And I think the short answer is pretty straightforward, and that is that the distance between what AI can do and what enterprise can actually operationalize has just become too large. So when we launched Microsoft Frontier Company, we built on the thesis that the forward deployed engineering model was largely the right one, that there's benefits of embedding engineering expertise directly into a customer environment, focusing on things like co-design, co-deploy, continuous improvement for AI systems at scale.

[03:04] And at the core of that, we really had to think about this through outcome driven by design. Because every system that we build at this point has to be able to deliver measurable business results. And I guess if I think about this, the best summary that I have is what I actually ask of my team, and that's to fall in love with our customer's problems and not our products.

[03:24] And maybe that means three things. Number one, how do we really help our customer at the operational layer of execution, really where workflow redesign, evaluation, observability are all important. The second is thinking about how do we make sure that we build reusable product capability by integrating all of the improvements like diags and workflow controls and audit trails and integrations that you really only see when you meet operations with your product.

[03:53] And then the third is the benefit back to Microsoft is how do we create industry patterns and repeatable use cases and workflows and integrations and implementation models that become scalable through the entire ecosystem? And I think about all of this as really focusing on the operational gap and closing it really being the work that needs to get done.

Danielle Rios:

[04:13] Yeah, I really like what you said there about being outcome-based focused versus these are our products that we're deploying. I think there's a really big distinction between those two things, especially when you're on the vendor side. Obviously you guys are trying to sell your products. And I

struggle with my own company and my own people, which I'm like, no, you got to start from what are we delivering to them that is of value? How do we measure that and making sure that we're delivering it? And yes, along the way we will use our products to do that, and they will sell themselves if you actually deliver the outcomes. And so I think that's super great.

Kevin Shatzkamer:

[04:46] I think that's spot on. I would even add to that a little bit. Because forward deployed engineering is a term that's gone through iteration over even the last 12 months. I think the term has really continued to morph to realize that it's not just building product, but it's really building operational capabilities and delivering outcomes at this point.

Danielle Rios:

[05:03] Yeah. And deeply understanding a customer's business. That's the key to it. I mean, you're on site usually, and you're working with real people in their jobs and understanding how their business works. And so you're leading the telco media vertical. Where is AI breaking down inside of telcos today?

Kevin Shatzkamer:

[05:18] Oh boy, we can talk forever about this.

Danielle Rios:

[05:20] It could be a whole podcast.

Kevin Shatzkamer:

[05:21] There's so many great demos out there and AI usually doesn't break down at the demo. It really breaks down at the seams of the operational environment.

Danielle Rios:

[05:30] Yep, totally.

Kevin Shatzkamer:

[05:30] And when I think about telcos, I think this is the single biggest transformation they'll ever go through. And I'm actually super optimistic that telecoms are going to capture this one. I think they have extraordinary technical depth, but they also have things like fragmented data. They have decades of accumulated systems, they have regulatory obligations, they have workflows that are really cross-domain across network, care, field service, security, and commercial operations.

[05:57] And when we think about all of this, a model can perform really well in isolation yet still fail when it doesn't have business contacts and permissions and observability and escalation scale. And for the first time, what I actually see is they're looking at this and they're saying, "This is structural

capability that I need to build.” And I think the industry is really converging around an AI-Native operating model thesis.

[06:19] And what's really cool is that a lot of the foundational capabilities are there. If I think about this through maybe three phases, phase one was this history of connecting people and things. Phase two has really been focused on how do I connect businesses and experiences and ecosystems? And really phase three, this AI-Native telco where I think about infrastructure that's sensing and deciding and acting and really continuously optimizing.

[06:47] I think phase two, they've spent the last decade doing. They've digitized the business, they've realized it's an imperative. And we talk about things like edge computing and digital identity and cybersecurity and IoT. And they've digitized and they built this platform. I think what makes this next phase super challenging and where they're really struggling is that the whole goal here of building a network that became a digital platform, it actually hasn't simplified the telco operating model. Instead, what it's actually done is it's created this digital services conglomerate operating on top of the network. And every time a new capability was added, it became another platform, another vendor, another API, an integration point, another data silo, another product, another life cycle journey for the customer.

[07:37] And all of this has just been this massive accumulation of technical debt. And I think right now where telcos are struggling is how do I simultaneously realize the technical debt burn down while also capturing this AI-Native opportunity? And I say it very simply, what we used to think about in the cloud world is migrate and modernize has become migrate, modernize, agentify.

Danielle Rios:

[08:00] Well, I mean it would've been super useful if they had made the transition to public cloud, let's say, over the last however many years I've been talking about it, because it would've set them up really nicely for AI.

Kevin Shatzkamer:

[08:11] For sure.

Danielle Rios:

[08:12] And so if they were late to that party, now they still have that work to do, and they got the AI stuff to do on top.

Kevin Shatzkamer: [08:19] 100%.

Danielle Rios: [08:20] And so we've had the move for 3G to 4G and 4G to 5G, now 5G to 6G, and the revenues haven't really followed the big investment of the new network. And so I think looking for ways to be more efficient in the network, AI is going to deliver in spades. And so talk to me about how it actually works at the Microsoft Frontier Company. Once your team is embedded, what happens?

Kevin Shatzkamer: [08:45] Yeah, we start simple. We start with an outcome, not a model. When I talk about phase three a little bit, what does that look like? It looks like the network becoming intelligent and customer service becoming intelligent. It looks like products becoming intelligent. So instead of the bundles of services, products become software plus AI plus connectivity plus data with outcomes.

[09:07] And I think it looks like the employee becoming augmented. Because AI productivity isn't just giving employees a copilot. It's really redesigning work itself around human and agent teams. And I think all of this is exciting, and there are growing proof points in the business functions across the board from AI customer care to AI coding to autonomous procurement, zero touch network operations.

[09:30] And I think we're at this point now where we've said the bottleneck is no longer an AI model, but I also think it's not cloud, it's not data, and it's probably not even engineering capacity at this point. The bottleneck is really transformation capacity. And I think that's our focus on what we want to bring to telcos. Every telco now has hundreds of potential use cases. They have dozens of pilots in play. They're doing multiple foundation models. They probably have thousands of agents, lots of data platforms. Obviously they're partially in the cloud from an infrastructure perspective, but how many of them can actually go fast from idea to production to measurable outcome and value to scaled capacity and capability? And I think that's the new problem.

[10:15] And I think it starts maybe with a couple of simple questions that probably aren't so simple to answer. Things like what is the business or operational result that you actually want to change, or what's the baseline, or who owns it, or what

constraints are non-negotiable? And I think they struggle with some of those questions.

[10:33] So the way we execute is really focused on observing the work. We don't just observe documented processes, but we actually look at how it works and how it actually runs. Where do we see people compensating for missing context? Where do approvals stall? Where do exceptions accumulate? And where does risk enter the system?

[10:52] And then from there we identify what is the smallest production workflow that we can actually focus on that proves this new operating model? Then we run it. And that's super important because a prototype is really good at telling you where something can work, but it doesn't tell you where it fails and it doesn't tell you whether an organization can actually trust it. And then once we're there, it's the learning loops. It's the learning loop for quality, for exceptions, for adoption, for business outcomes, and for human in the loop because our objective isn't really this clever agent. It's really some kind of agentic harness that's repeatable and lets our customers build, govern, observe, and improve lots of agents over time.

Danielle Rios:

[11:31] Yeah. And what's interesting is that it's not just an engineering effort, this isn't just building a product or a solution for a particular customer. But there are a lot of human resource focused motions, studying the way people work, redesigning their jobs. That's going to require, like you said, transformation. And it's not just a matter of writing a new tool and deploying it and say, good luck guys. You really have to bring the humans along and tell that story of why this is not only good for the company, but also good for them as individuals.

Kevin Shatzkamer:

[11:59] Yeah, I think about that, DR. And I think what's interesting is when we've thought about transformation to date in telco, we've thought about it as these big heavyweight programs that last years. And I think what we're trying to instill here is a whole new discipline that it's not milestone driven and it's not these long-term engagements, but it's really small capability driven.

Danielle Rios:

[12:21] Totally.

Kevin Shatzkamer:

[12:22] And when I think about transformation, it's how do we actually snap it to what engineering cadence looks like? And we

know engineering cadence at this point looks like two week sprints. So how do we think about capability transfer in a continuous way, how do we think about change management in a continuous way. And have it snap to that two-week cadence that we try to instill from an engineering discipline perspective?

Danielle Rios:

[12:43] Yeah, I mean, it's exactly how Totogi works. And we call them capabilities as well, and we have four deployed engineers as well. And so when we're working with a customer, we have two meetings a week where we're showing new advancements because that's how fast AI allows you to move. And so yes, it ends up being these little sprints where you're like, "Okay, you gave me feedback on this. We found some gaps. We revved it. We'll see you on Thursday. We'll see you on Monday. We'll see you on Thursday."

[13:09] And that ability to show results so quickly allows customers to learn as we're building. And so we try to do almost twice weekly. And so telcos know they can't outbid Microsoft or other hyperscalers for AI talent, some of the most expensive talent right now in tech. And so do your engineers train the telco teams to eventually take over their own work, or is really Microsoft becoming their AI team?

Kevin Shatzkamer:

[13:36] Yeah, maybe that's a false choice. I think if Microsoft becomes our customer's permanent AI department, we haven't built a durable operating model for the industry. And our goal here is really transformative change at the industry level. But also I think pretending that every enterprise should independently recreate a hyperscaler or stumble through the problems that already have solutions is probably not productive.

Danielle Rios:

[13:59] I agree.

Kevin Shatzkamer:

[14:00] So I don't think this is about really taking over telco operations. We've tried that before as an industry, and I think most telcos have re-insourced a lot of their operations. It's really how do we bring tools and technology and methodology and process, but also expertise to code with, design with, as a motion because we're looking to embed our teams for a period of time. And when we leave, we leave behind a structural capability.

[14:25] And I think the capability interlocked between telco and Microsoft is pretty clear. We'll bring platform depth and we'll

bring engineering patterns and product knowledge, and also what we're learning across other operating environments. And the telco contributes something we can't manufacture, the network context, the customer context, the regulatory judgment, the operational excellence, and the accountability for the business outcome.

[14:49] So our job is really a build with motion from the beginning. And that means shared architecture, shared decision, shared governance and deliberate transfer of knowledge. But even more importantly, it's really shared outcomes. My team is not successful when we've delivered. We're successful when the outcome that we intended to create for the telco is realized.

Danielle Rios:

[15:09] Yeah, that's awesome. And so we've been talking about how AI really is defining every job, not just in telco, but just around the world. You and I both have kids that are going off to college for the first time. I have a daughter who's about to be a freshman at UC San Diego, and I have a son who's a junior at Rice. You just dropped off your two twin boys at two different universities. And so what are you telling them about the impact of AI on their future as they are entering the workforce in the next four years or so?

Kevin Shatzkamer:

[15:41] Yeah, I'm probably not telling them at this point to try to predict which jobs AI will replace or will change because the work inside just about every single job is going to change, as you said. And what I'm telling them is to build judgment, to learn how to frame a problem, to ask a better question, to think about system dynamics, how to test and answer, how to take accountability for outcomes, because I think those are the transferable skills that will matter.

[16:06] Their generation is really going to enter the workforce with a set of capabilities that we didn't have when we entered the workforce. We are really at the rise of the internet. And I remember I was working for Cisco at the time and John Chambers said, "The internet's not just going to change what people learn, but it's going to change how people learn." Because before the internet, there was value in holding and retaining knowledge. But over time, what we've seen with the growth of the internet is that knowledge is at your fingertips. And I think in this next world, it's not just knowledge, but actually task execution that's at your fingertips. And when we think about that world, the problem probably belongs less to

the person who has every answer and more to the person who can frame the right problem.

Danielle Rios:

[16:47] That's super interesting. I've been telling my kids, build expertise. It doesn't matter what it's in, pick a topic and be an expert. I think with AI just summarizing everything, it's really easy to just be lazy about it and have it just summarize the book or summarize the topic. And I think, just like what you're helping telcos do, documentation of expertise is really what's going to drive AI. And if you can do that and be very good at that, you will be the master of AI and you will have a great future.

Kevin Shatzkamer:

[17:17] 100%.

Danielle Rios:

[17:18] Well, Kevin, this is such a great conversation about what's going on in Microsoft and the Frontier Company and the FDEs. I'm so excited about what you guys are doing for telco. So great talking to you. Thanks for coming on the podcast.

Kevin Shatzkamer:

[17:29] It's my pleasure. I appreciate the opportunity to chat.

Danielle Rios:

[17:32] Awesome. Stick around. We end each podcast with a "Telco in 20" takeaway. I've got two minutes to tell you something you need to know. Kevin's diagnosis is right. Transformation capacity is the bottleneck, but there's a question underneath it. Transformation capacity for what? Telcos have a pattern. Every technology shift sends them hunting for a new revenue story. The killer app, the new business model, the thing that will finally make the network pay for itself. And every time they miss what's actually sitting in front of them.

[18:15] AI is not a new revenue play. It's the chance to rebuild the telco operating model. That's the opportunity. And the operators that capture it won't be the ones with the most pilots or the biggest AI budget. They'll be the ones that rewired how change gets done. Sprints measured in days instead of two-year programs, shipping in weeks, not quarters, fewer meetings about AI and more AI in production.

[18:42] So go ask your team one question, "What's our cycle time from idea to production?" If the answer is months, you don't have an AI problem, you have a transformation capacity problem. And that's one worth solving. Want to ship in days instead of months? DM me on LinkedIn or X @TelcoDR, and I'll show you what building with Totogi looks like. Or come find me

Episode 149 | Microsoft deploys the FDEs (Kevin Shatzkamer)
Released September 15, 2026

at TM Forum's Innovate Americas on October 6th in Dallas.
Yee-ha. Ryan Alvarez, CEO at VIVA Bolivia, will be joining me on stage to talk about how he's using the Totogi Ontology to run every cell site as its own P&L. Until then, tune into more "Telco in 20" episodes. Like and follow and leave us a five-star review. Don't forget to sign up for my awesome email newsletter on TelcoDR.com, and check out our killer YouTube channel and hit that subscribe button. Later, nerds.