

Episode 135 | Telco has a context problem (John Abraham)
Released March 3, 2026

- John Abraham: [00:00] The challenge that I anticipate if we don't have an effective way to solve the context problem is that we risk ending up with a number of agentic silos.
- DR: [00:10] Yeah, and I think there's something really big that's missing from all these approaches. Everyone is tackling it as a data problem, and it is a data problem, but you're completely missing a very important piece, which is your business processes and rules, how you actually make decisions.
- Announcer: [00:33] This is "Telco in 20," a podcast that helps telco execs achieve a competitive advantage with AI and the public cloud. It is hosted by Danielle Rios, also known as DR. Today we're talking to John Abraham, partner and principal analyst at Appledore Research.
- DR: [00:52] Hi, guys. I'm DR. Last month, OpenAI launched Frontier to help enterprises deploy AI agents, while Amdocs launched AOS, an agentic operating system. Both are focused on making agents run; neither is focused on making agents think. The hard part, the part nobody is talking about, is how to capture the way your business actually makes decisions. That logic currently lives in two places: buried in vendor code that nobody fully understands and inside people's heads. But this is the layer AI needs: what's allowed in your telco operations and what is not. That's not in your data lake, that's not in an API. And no agentic system, no matter how shiny, can make decisions if it doesn't have that. When you create a layer that captures how your business thinks, that layer compounds in value. Every decision makes it smarter, and that knowledge is owned by you, not your vendor. Remember, Amdocs is the king of the \$50,000,000,000 services industry in telco. They profit from keeping that knowledge locked in their system and away from you.
- [02:01] Today I'm talking with John Abraham, partner and principal analyst at Appledore Research. We're going to dig into why this context problem is a real blocker for AI at scale, how the Totogi ontology gives AI the information it needs to actually work, and why your context layer is your true competitive moat. So, let's take 20. John Abraham is partner and principal analyst at Appledore Research. Hi, John. Welcome to "Telco in 20."
- John Abraham: [02:30] Hey, DR. Great to be back. Thank you for having me.

DR: [02:33] It's always great to talk to you about what's going on in the industry. And so you're leading the digital enablement practice at Appledore, and what's so great about you is that you have the opportunity to talk to both vendors and CSPs. And so what are the big topics people are asking you about?

John Abraham: [02:51] Well, I'll give you the bird's-eye view from a digital enablement perspective. First is AI. It's everywhere right now. But one of the notable shifts is that people are talking about the underpinnings of AI almost as much as the flashy application. The way I see it is it's a three-tiered pyramid with a data layer at the bottom, the ontology or the semantic layer in the middle, and the application layer at the top. And what we have seen in the past three, four years has been a lot of emphasis on the application layer. What is a new fancy thing that AI can do? And I think increasingly we are seeing that shift towards the bottom layer because that's what will actually enable and unlock the value of AI for telcos.

[03:35] Second, enterprise. That remains a big frontier for telco. They want to be considered on par with some of the hyperscalers in terms of going up the value chain because a lot of them are partially there already because they provide the underlying connectivity, and they want to go up that chain.

DR: [03:52] Absolutely.

John Abraham: [03:52] And the third is monetization. And this is an interesting one because the entire emphasis is on two things. One is to make it more agile, and the other is to make it cheaper because telco monetization is notoriously expensive. It's got a lot of legacy in there, and it is not agile or dynamic enough. So those are the twin problems that are being addressed within that domain.

DR: [04:12] Yeah, absolutely. And I totally agree. I think these are going to be super big topics at MWC, especially AI, and I think agentic AI. And so I recently had Rick Lievano from Microsoft on the podcast, and he shared with me this really interesting story about an operator, AT&T, in the United States that had set up that top layer that you were just talking about, AI agents, and they were getting ready to deploy them to production and they stopped because they're like, "Are we really in control? Do we really know what these agents are going to do?"

[04:42] And so what are the problems that you see telcos talking about when it comes to deploying AI at scale?

John Abraham: [04:48] Yeah. So I'll capture them on what I call the three Cs, the main challenges. And the first is obviously cost. It's got multiple dimensions to it. On one side, that is the actual issue of the money it costs. There's a lot of hidden costs associated with it, some of which many telcos are just beginning to uncover. And the other side of it is about what should we be expecting out of spending all this money? What is the value in it?

DR: [05:11] Absolutely.

John Abraham: [05:11] I know of a Tier-1 telco who's using up roughly 5,000,000 tokens a day, and they don't know where they are right now. Is it just the early stage, or this is just a high, and we can see better efficiency in the future? So these are still open questions.

[05:26] The second C is about context. And this is a big one because a lot of the times agents are operating at least partially blind. They don't have sufficient data or context that helps them guide to the next action. And we don't have clear ways where telcos have solved it yet.

[05:42] And the third is communication. And that is: how do you enable a friction-free mechanism for agents to communicate with each other, keeping in mind that they need to be coordinated, they need to understand each other's intent?

[05:56] So these are the three big issues that I would call out that's actually holding back that scaling, especially within telco environments.

DR: [06:03] Yeah, I was looking up the cost of 1,000,000 tokens over the course of the last few years, and I think even two years ago, 1,000,000 tokens was about \$30 and we're now down to 10 cents.

John Abraham: [06:16] Yeah.

DR: [06:17] And I think in the wars with LLMs, the prices are going to continue to drop.

John Abraham: [06:21] Yes.

Episode 135 | Telco has a context problem (John Abraham)
Released March 3, 2026

DR: [06:21] And so I think telcos would be happy to spend a lot of money, millions of dollars, provided that the value was there. And I think that's the key.

John Abraham: [06:32] That's right.

DR: [06:33] And I think to really get value, you need to go all the way. I think people are dabbling in productivity improvements. At the end of the day, you still need a person reviewing what AI is outputting.

John Abraham: [06:43] Yes.

DR: [06:44] I think really to get that value, you're going to have to do some bigger, scarier things. But I think the second thing that you said, context is going to be very important.

John Abraham: [06:52] Can I just add one point to what you just said, DR? I fully agree with you. I think what's unique about the AI wave is that telcos don't have the option to sit out while the technology matures, which they have often done with previous generations of different types of technology, where they would wait for a little bit of time so that it matures before they try and embrace it.

DR: [07:13] Yeah. Well, AI's felt like a speeding train. You can't sit on the sidelines and wait for this to settle and see who's going to win. You have to jump on the train. If you wait, it's only getting faster.

John Abraham: [07:23] That's right.

DR: [07:24] Yeah. So I think something that you've been saying is this management of context. How do you see telcos tackling this problem and starting to manage context for all of their AI projects?

John Abraham: [07:36] So that's actually a tricky question because it's quite clunky. That's a short answer to that question, DR.

DR: [07:43] No plan is what I think it is.

John Abraham: [07:44] There is no clear path to it yet. I think it's fair to say that we are seeing numerous attempts at trying to solve the problem, but the challenge is going back to the previous point

you said. It takes time, and the entire AI train is moving so fast that I don't think we have enough time to spend on building up the context before we get ready for AI.

[08:05] I'll just talk through what I see, a couple of approaches. Standardized APIs is one of them. TMF has their own standards. GSMA has some standards. So this is not widespread yet in terms of being able to address context for AI, but there is some activities going on there. Some operators are trying to extend the MCP, Model Context Protocols. And again, I understand there are some proposals in front of TMF right now to bring some sort of a standardization to it. And as we know, these are things that take time to get a consensus across both the vendor and operator landscape on something as critical as this.

[08:41] There is TMF's Model-as-a-Service or MODaaS, which is mainly about standardizing how AI/ML models are exposed and consumed as a service, which means that it makes it easier for agents to call the same set of tools. Again, very early days. I think they just have the framework for that getting ready.

[08:59] So the challenge that I anticipate if we don't have an effective way to solve the context problem is that we risk ending up with a number of agentic silos that just exist in and of itself and doesn't interact with each other, which I fear may not be too different from the siloed landscape that many telcos have today.

DR:

[09:19] Yeah, and I think there's something really big that's missing from all these approaches. Everyone is tackling it as a data problem, and it is a data problem. You do need to get your data to the agents, absolutely. But you're completely missing a very important piece, which is your business processes and rules, how you actually make decisions. APIs can't pass that. MCP doesn't pass that. Some people are building a data lake. Doesn't do it.

[09:49] And so, everyone's focusing on the data part, which you do need to solve that. You need semantic consistency, you need access to the data. The whole part about context is the description of how all of these entities and objects relate to each other, what they're allowed to do, what are their options for decisioning.

[10:10] And so I think this is where telcos absolutely need to spend time. This is the investment layer. I wouldn't waste time building a model or fine-tuning a model. Models are commodities. The minute a new one comes out, everyone gets access to it. Everyone's going to use the best one. Microsoft, Amazon, Google, they're going to come out with developer tools that help us build AI workflows much more quickly. But that context layer, that's 100% your telco. This is the piece that makes AI valuable. I sometimes think of it as their moat. It's something you should absolutely be thinking about and not doing it as an afterthought.

[10:48] If you don't get your context, you're going to have a real big problem. You're not going to be able to scale your AI. And if you don't solve your semantic chaos problem and how you codify your business rules, this is the meat of AI. And I don't really see a lot of telcos or even vendors, except for Totogi, talking about that, mainly because we've been living it.

[11:10] And so you and I worked together on a white paper about this recently on ontologies, and so you and I have been talking about it a lot and thinking about it. How do you think an ontology can start to solve this context problem for telcos?

John Abraham:

[11:23] I think one of the key takeaways I had when I worked on this paper is that we don't have a consensus on what ontology means. It means different things to different people.

DR:

[11:32] Mm-hmm.

John Abraham:

[11:33] For many people, ontology is just the semantic backbone that can map and synchronize data across different systems, which by itself is quite a big deal.

DR:

[11:44] Yep, very important.

John Abraham:

[11:45] Absolutely. Especially in an operating framework like telco with multiple multivendor legacy systems, it can be quite useful.

DR:

[11:52] Mm-hmm.

John Abraham:

[11:53] So one of the more common definitions of ontology is the way Palantir frames it, and there are others. I just picked them up because there seems to be at least reasonably

accepted among a number of different uses, and they define three different components that make up ontology. The first is the semantic layer, and this is how things are defined and how they relate to each other. And you can call them nouns.

DR: [12:18] Mm-hmm. The nouns of your business.

John Abraham: [12:19] Exactly. The second is the kinetic layer, which is an understanding of how things work, how they interact with each other, how they behave, in other words, the words of the business.

DR: [12:28] Mm-hmm.

John Abraham: [12:29] And the third is the dynamic layer, which is about how to make a decision and how to take an action. So these three bits together constitute ontology. I think that's a crucial distinction that many people seem to miss.

DR: [12:44] Yeah, I think you've captured it. It's a big word that people have heard before, but they don't really know what it means. And I think you're totally right. People sometimes use it to describe, "Oh, I have semantic consistency like you would achieve in a data lake," but a data lake can't make decisions. Some people refer to an ontology as a knowledge graph, but again, a knowledge graph can't take action.

[13:04] And so the way Totogi is using the word ontology as, it has semantic consistency, it has the knowledge graph of what your business can and can't do. But that last really cool part, it can make decisions and take action.

[13:17] And what's so beautiful about it is that it compounds. This isn't just an integration or an API that you write and you use over there. As you add to the ontology and improve the decisioning or change the decisioning, everything else can use that going further. So any AI that uses the ontology benefits from that.

[13:37] And so the agents aren't unique things that you build and deploy, but rather they're part of this living fabric. I like to describe the ontology as it's the layer where data becomes decisions and decisions become actions.

Episode 135 | Telco has a context problem (John Abraham)
Released March 3, 2026

John Abraham: [13:50] Yeah. I think that point you mentioned about compounding, that's quite significant. You start small, and then you keep on adding to it. And oftentimes, at least in the telecoms industry, we are not used to that level of progressive improvements. We have very marginal improvements over time, but I think ontology changes that in a completely radical way that we haven't seen in telco before.

DR: [14:14] And the beautiful part of it is that you don't have to model the entire telco. You can take a small subset, and we've seen people start with these little small projects, but then they just start adding more things. We started with a small problem solving dormant cells at a lower part of the organization in telco, and they're like, "Oh, our cells fall asleep, and we don't really know about it."

John Abraham: [14:33] Yeah.

DR: [14:34] Okay, great. It's like \$1,000,000 or \$2,000,000 problem to fix. It has value, as we've been talking about, but it's not a huge value. But then we started connecting different things to it, and now we're at the CXO level, and they're like, "Can you connect our fuel information so that we can make decisions on how we're deploying fuel across our country to invest in our different cells?"

[14:55] Now we're getting into millions of dollars of impact, but it all started from a very small area. We haven't modeled the entire telco, and they're getting value, and I love that about Totogi. I love that about our product. It's pretty amazing.

John Abraham: [15:09] Yep. One of the interesting things that we spoke about some time ago, DR, was how you are framing this as operators don't need to have the data lake in place before they start work on ontology.

DR: [15:23] Mm-hmm.

John Abraham: [15:23] They can start with whatever they have because there's a lot of money and time and attention being spent on trying to have that perfect data lake.

DR: [15:32] You're exactly right. Let's just fast-forward. Let's spend \$20,000,000 and snap our fingers. It's two years later, and you got a data lake. Congratulations. Your data's all clean.

Meanwhile, on your normal systems, business is going on, you're still using them. Now that you have this clean data and you have some insight that you want to take, you're going to add AI to it, you need to take action. Where does that happen again? Does that happen in your data lake, or do you have to now tell people over in the systems of action that are working to take that action?

[15:58] You need to put your semantic consistency on your live systems. A data lake is a side quest. It's a \$20,000,000 multi-year side quest. Stay on the AI path. Build an ontology, invest in your context, make that work, and it will pay off massively with your AI. That's how you get value. And AI's making decisions? Oh, my gosh. Now you can scale. Now you can build out all the AI things you want to do.

[16:23] And so I'm obviously very passionate about it, very excited about what's going on, and really excited to see everyone at MWC. But I think what's interesting about you, John, is I haven't seen you in a couple of months in person, and I noticed you have grown a beard, and so I've never had to maintain a beard personally. I hope I never have to, but we've been talking about maintaining context. How do you maintain your beard?

John Abraham:

[16:50] Well, I'm just six months old in this, maybe slightly less than that. It's fascinating, every time I get on a call, people pause for a second trying to figure out who the heck is this person before recognition dawns.

DR:

[17:02] Yeah, that's funny.

John Abraham:

[17:03] So MWC, that's probably going to be the case. One of the things I found out about this beard thing is that it's incredible the number of accessories and types of creams and oils that you can get out there. I was completely oblivious of this. You can get oil, shampoo, conditioner, balms, different kinds of stuff. I think the brand I would recommend is something called Viking Revolution or something. That's quite a good one. So yeah

DR:

[17:26] Awesome. Well, I don't know this world. I'm not surprised as a woman. The creams and the shampoos and the treatments, I can only imagine. But it looks great on you, and I can't wait to see you in person at MWC.

Episode 135 | Telco has a context problem (John Abraham)
Released March 3, 2026

John Abraham:

[17:39] Thank you. We are looking forward to it.

DR:

[17:41] Yay. 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. John mentioned costs as one of the things holding telcos back from AI at scale, but let's be real. Token prices have dropped from \$30 to 10 cents in two years. Cost isn't the problem. Getting value from AI is. So how do you actually get value? It's easy. Tie the results to a number your CEO cares about, increasing revenue or decreasing cost. That's what we did at Totogi with a Tier-1 operator in the Middle East. We started small, working on identifying dormant cells, which is a \$1,000,000 revenue loss problem. AI identified the issue, the ontology knew what actions were valid, and the system executed. No tickets, no human intervention.

[18:40] But here's what mattered. Every decision made the ontology smarter, so the next problem was easier to solve. They connected more systems, then more. Same foundation, compounding in value. Now we're at the CXO level, making fuel deployment decisions worth tens of millions of dollars, all built on top of that first use case. True value delivered from AI.

[19:04] We're dropping this episode during MWC. If you're in Barcelona, head over to Hall Two, booth 2G51 to see Totogi in action. Shoot me a DM on LinkedIn or X @TelcoDR, and we'll set up a time to meet. Until then, tune to more “Telco in 20” episodes, like and follow, and leave us a five-star review. Don't forget to sign up for my no-BS email newsletter on TelcoDR.com and check out our awesome YouTube channel and hit that subscribe button. Later, nerds.