

DR: [00:00] I'm DR and this is Telco in 20.

[00:14] Microsoft is leading the generative AI race among the big three hyperscalers with its huge investment in OpenAI. CEO Satya Nadella continues to make major moves in this space, like acquiring startup Inflection AI for \$650 million and announcing plans to build a \$100 million data center that includes an AI supercomputer.

[00:32] But while a lot of organizations are experimenting with AI, it's time to start integrating AI into operations to drive transformative change. For telcos, we have a long road ahead of us to really operationalize AI workloads. Besides getting our people trained in ready to use AI tools, we need to get our data cleaned and into the right databases so AI can work its magic.

[00:53] And who better to talk about this than Jason Hogg from the AIOps team at Microsoft, he's been working with telcos around the world to bring AI into their organizations in meaningful ways. Today we're going to dig into where telcos are on their AI journey, how Microsoft is helping operators get their data ready for AI workloads, and how you might need to rethink your organizational design to optimally use AI. So let's take 20.

[01:20] Jason Hogg is a general manager for AIOps within Azure for operators at Microsoft. Hi Jason. Welcome to Telco in 20.

Jason: [01:27] Hey Danielle, good to see you again.

DR: [01:29] I'm so excited to talk to you about AI, what Microsoft is doing. You guys are like the belle of the ball, at least in the AI world, so can't wait to talk to you.

Jason: [01:39] It's amazing. There's just so much going on. It's a fun period to be working in technology, right?

DR: [01:42] I agree. And so, we were on a TelecomTV, AI-Native Telco panel together last year. And on that panel and since then, there's been a lot of discussion about where AI workloads will reside, whether they'll be on-prem or in the public cloud. And so, coming from Microsoft, what's your take on where people will be running AI workloads?

Jason: [02:05] I don't think it's an either-or question. I think a lot of it comes down to what the requirements of the workload are, plus

what the location and availability of the model are. I think there's lots of really great examples of techniques like clustering, forecasting, anomaly detection that run in near real-time or in batch scenarios. The public cloud is a perfect location to be running that sort of a workload. I think there's other examples like the RAN for example. There really is a need for real-time inferencing. And so I think those workloads will likely run at the far edge. And then, of course, AI within the handsets. We're starting to actually see Samsung handsets and so forth with GPUs inside them that allow applications to start doing inferencing on the far edge. And that's super interesting as well.

DR: [02:45] Yeah. And so, when I start to think about telco workloads and just the vast data that telcos have, I think today the data is on prem and most of this data is trapped in vendor databases spread literally all over the organization. And if you're a big group, Vodafone group, or MTN, not only is it spread across your headquarters, it's spread across opcos in different countries. And so I believe that your AI results will only be as good as your data. I think if you are training garbage in, you'll get garbage out. And so what is Microsoft doing to help telcos manage their data and get it ready for AI workloads?

Jason: [03:28] Yeah, that's such a great question because that's been one of my observations is these sort of trapped siloed data sets that are really sort of limiting the ability to actually take full advantage of their data and the insights that they've got. And it's funny, it reminds me of 11 years ago when we first sort of started Azure, each of the services within Azure had its own means for emitting telemetry and its own lifecycle management technologies.

[03:49] And what it meant was it was very hard to actually get a full understanding of Azure as a single service. And so that's where we invest significantly in an area that's called AIOps, which is really about two things—Basically, the ability to actually ingest and analyze massive quantities of data from different sources. So you could actually really understand systematically what was going on across the services. And then in lifecycle management, you could actually do complex deployments. And so, what we're trying to do is actually help take the same technologies that we use to actually mature Azure and take it to telco so they can actually start removing these barriers, moving massive quantities of data cost-effectively to the cloud, and then sort of running AI, not just on individual data sets, but on the

combination of data sets, the data mesh if you will. So that's the big part of what my team's focusing on.

DR: [04:34] Are you trying to ingest a ton of data? Are you moving data up into Azure? Is that a big part of it or are you leaving the data where it is and cleaning it for telcos and organizations?

Jason: [04:45] Yeah, that's one of the challenges. There's not really another industry like telcos where the actual sheer quantity of data that sort of gets emitted from these networks. So in some cases, we're able to take the data right from the network functions, move it to the cloud, run AI on it there. In other cases, it's absolutely cleansing that we do at the edge, removing PII, maybe even do some local aggregations, those sorts of things. So yeah, guess both would be the answer.

DR: [05:08] A little bit of both. Yeah. So hands down, as a hyperscaler and as a company, Microsoft has been kicking everyone's butt and leading on the GenAI front. It's been incredible to watch what your CEO Satya has been doing. And so what are some of the Microsoft AI tools and offerings that telcos are beginning to adopt and leverage either internally or for their networks or otherwise?

Jason: [05:30] Yeah, Azure OpenAI is how we're actually enabling operators to get access to the underlying generative AI that powers ChatGPT. One of the technologies I mentioned in our AIOps suite is called Azure Operator Insights, and that's actually the platform that we've built specifically for telcos for ingesting, analyzing data. We've actually got generative AI integrated into that along with a RAG implementation with the goal being initially to help operations teams start embracing generative AI.

[05:55] So we can actually take an alert that gets raised from network data and then actually synthesize recommendations for how to go about mitigating the alert using telco-specific technical documentation, support docs, telco-specific RCAs, those sorts of things. And basically starts increasing the efficiency of operations teams.

DR: [06:13] Yeah. The hard part has been the human interaction and taking the work away from the humans and putting it into AI and the recommendations coming out and then applying them, either using the recommendations or just letting the AI flow through. Are people fighting that? Are you seeing any HR

struggles with letting the results actually get implemented so that it's not just, "Oh, that's interesting information. Okay, I'm going to put that to the side and I'm going to go ahead and do my job the way I was doing it before."

Jason: [06:43] No. What we're seeing is actually a strong desire from the people using the tools to actually be able to work more efficiently. One of the studies that I sourced showed that operations teams spend up 20% of the time just actually searching for solutions to problems. One of the other examples of how AI is being applied, that was super exciting at Mobile World Congress, there's a ton of work we're doing around voice AI now where we can have agents participate in phone calls and actually start to look for signals like fraudulent scenarios where somebody may be trying to actually encourage your parents to hand over their passwords, their credit card numbers, or access to their computers, and the agent can now jump into the call, warn the person on the call.

DR: [07:20] Oh, wow.

Jason: [07:21] Yeah. This kind of stuff is super exciting and it's essentially breathing life back into those traditional voice-based workloads. That was exciting.

DR: [07:29] Yeah. There was a lot of interesting shenanigans, I guess that happened last November with OpenAI and board turmoil and Sam Altman, is he in, is he out? And certainly Satya had to get involved I think at some point. And so did you guys see customers start to back off in terms of, "Hey, maybe we should pick more than one AI vendor?" Are you guys getting a lot of questions around that?

Jason: [07:55] It's funny. It hasn't come up in any of the conversations that I've had. I think it was sort of a point in time anomaly. If anything, it probably generated more attention around OpenAI and the underlying models. Telco is a very mature industry, so I think in most cases they recognize it's really about what's the problem that you're solving and how fast can you get it to market and does it affect my availability or my ARPU or COGS in any way? So those are more of the actual discussions we're having, which is how do we actually use these technologies to change the operator's business if you will.

- DR: [08:23] It's shocking they're not worried about vendor lock-in with their AI vendor. They're so concerned about vendor lock-in on their cloud provider. And so I think they'd be asking questions around cloud and maybe they're not that mature with their AI workloads.
- Jason: [08:34] I think the operators are looking at lots of different solutions to lots of different problems. So obviously OpenAI is not the only technology in the market, and I think many of the operators are looking at different stacks. I think in the generative AI space, probably what I sense the most is that there's lots of R&D teams looking at lots of different technologies. And probably a fear that it's not really where to place their bets, which is again, why I think it's sort of important to up level and focus on the scenarios and the requirements, and then the underlying technology will change.
- [09:04] And it's sort of 4.5, is it 5? Are you using your own RAG implementation? What vector database are you using under the covers? All of that sort of stuff shouldn't be interesting to a telco, in my opinion. It's more about what business problems you're solving, and that's I think where the discussion needs to actually move to.
- DR: [09:17] Yeah. So let's talk about that. So I do agree it's still pretty early days and-
- Jason: [09:21] Correct.
- DR: [09:22] I saw an interesting quote from the Adobe investor call of we need to move from experimentation and start getting things into production.
- Jason: [09:30] Oh, that's such a great quote.
- DR: [09:32] It was awesome. And I see that even with my own team. They're like, "Oh, I did this cool thing, but then I put it off to the side and went back to working the way I was working," and I keep telling them, "You guys need to ship shit. Just ship it. Just get it into our workflow and get it done. And so we're using it in customer support tickets. I think there's a huge volume of solved tickets that's great for training to create a RAG vector database on. And there's lots of formulations and examples, but where else are people starting to really put AI into production and using it within the organization?"

Jason: [10:03] I think that Adobe quote is such an important one because as you walked around Mobile World Congress, it was like every single stand, every single vendor had AI written somewhere on the dashboard. And I think the challenge with AI, which is very different from many other pieces of software, is that it's very easy to actually build a prototype. It's very easy to build something that looks good in one specific use case, but generalizing that and putting it into a production environment, completely different.

[10:26] I'll just give you one example of that. I mentioned we've got this platform, Azure Operator Insights, which includes ability to ingest and analyze data using sort of traditional machine learning techniques, clustering, anomaly detection, those sorts of things. And we're in the middle of integrating generative AI into the platform. But Microsoft takes the actual testing and the actual release of generative AI and AI in general, just extremely seriously. So it's taking me probably six months longer than it would actually to release anything using generative AI than it would use traditional techniques to ensure that basically, we've thought all about the corner cases, the edge cases, the risks of using the technologies. It's very easy to create a prototype, very different to actually build something that works in a production environment.

DR: [11:06] Our approach has been to start with an easy case and nail that and then tackle the next easiest case. And then see what's falling out, what AI can't solve. Just make a list and just keep tackling it. And I wrote a blog about this. We were doing this with our support team. We started with the idea of, "Could AI solve a hundred percent of our tickets?"

[11:25] And at first it was a very small percentage, but now we're up to 57, 58, and what's the next-

Jason: [11:32] No, I love it-

DR: [11:32] Most common easiest thing for us to go do? And we put it into production and you have the corner cases and the problems that are falling out. We have a human watching over it to see if it's okay, but it's been a very iterative process over the course of months. It's not quick—one and done.

Jason: [11:48] I love it. And that's exactly the same way that we build models from managing the backend of Azure. We've got this

experimental phase where we sort of specify a hypothesis, test it, work out what the accuracy is and how we can improve it either with more data or additional types of data. So that's exactly what we do on the back end. And I think then with telcos in particular, one of my observations has been that need to actually evolve and continue to mature the technology is a new paradigm for them as is the need to actually get access to more data and different types of data very quickly. Quite often, like you mentioned, siloed data sets makes it tricky to actually get hold of some of the data sometimes.

DR:

[12:21] Yeah. One insight that we had, and I don't know if you guys have been seeing this, telcos have tons of people, thousands and thousands of people in their organizations. And to date, our organizations have been designed around people. Managers, teams, tasks, like for example, in support, they move from level zero to level one to higher levels of pay, higher levels of complexity, higher levels of technical competence.

[12:47] But one of the breakthrough ideas that we had with our support team was we're applying AI to a human-centric design org. We need to flip it upside down. We need to think about what AI can do and then have humans that are shepherds of the AI. They're either improving it, they're tooling it, they're watching over it. And that really helped us kind of have a breakthrough of let's redesign the organization to be AI-first. And so, as you guys work with telcos, are you guys seeing them have those breakthrough thoughts of let's redesign the org?

Jason:

[13:16] Absolutely. I think the operators that are making the most effective use of data and AI are thinking more about what the outcome is and then how to get there with the combination of the people and the tools that they've got. You sort of mentioned the operations or the customer support space. I think there's just a massive opportunity there to actually start thinking about those people's roles as not just solving the tickets, but as curating the knowledge base that's actually powering the underlying model.

[13:39] There's a major telco in the UK we're working with right now where they're restructuring their organization with people that have sort of this growth mindset approach to actually thinking about not how to automate existing process, but what's the right process with the technologies you've got to automate it. There's a really good example of this in the deployment

spaces. Right now, quite often, software is versioned and deployed sequentially.

[14:02] If you could actually deploy multiple network functions in parallel with the more modern deployment techniques, then you can actually get new functionality out much faster and more predictably. So it's just a different mindset and the kind of mindset change that we had to embrace when we were evolving Azure from a collection of services to one single service.

DR: [14:19] Yeah. No, I totally agree. I think you just have to almost flip it upside down and think so differently. And so can you guys share examples of telco customers? Who are some of the leaders that are emerging and some of the customer case stories that you guys have of people applying AI and getting some really awesome business value and results?

Jason: [14:37] I think there's lots of sort of public examples in the United States. We've got an amazing relationship with AT&T, the underlying NFVI layer, their network cloud infrastructure. We've evolved, that's become our Nexus platform. Nexus is designed with the ability to actually emit telemetry that we can run AI on from day zero. So that was just one of our critical learnings we took from Azure integrated there. Got an amazing relationship with Three UK, and there's actually a really great video on YouTube if anybody wants to look for it.

DR: [15:03] Yeah, we'll link it in the show notes.

Jason: [15:05] Yeah. Okay, cool. That'd be awesome. They're basically talking about how they've used Azure Operator Insights to actually break down the silo, generate new insights, and they're one of our earlier users of the Copilot and AOI capability. You sort of talked about the relationship between people and AI, and one of the things I want to emphasize is we don't really look at AI as necessarily removing the need for people if you will. We've got this incredible brand with Copilot and Copilot is all about really thinking of AI as an assistant for people.

[15:30] I don't know if you've seen Copilot for GitHub. It's an incredible tool that allows code to be synthesized for developers, but it's not meant to replace developers. It's meant to actually make developers more productive. At the end of the day, they're still responsible for checking in the code. It's not that the code's writing itself and you no longer need developers.

So I think it's just an important thing to keep in mind. This relationship between people and AI and the role that people play is still very much important and critical.

DR: [15:55] Yeah, I think it's going to change work just like with any technology advancement. Maybe the old jobs that you used to have may go away or change dramatically, but I think it'll create a whole set of new jobs. We're hiring tons of people who understand how AI works. Prompt engineers, I can't hire enough of them.

[16:16] And so, maybe five years ago they would've been software developers, but now we're looking for people who really understand how LLMs work and how to write amazing prompts. And so I think just work is changing. It just is. Now you have a Copilot that maybe is assisting you, but I think it's going to just up level, and if you want to keep making more and more money, I think the people who understand how AI works, that's where the money's going to be.

[16:41] And so, using AI is like climbing a mountain. Every day I wake up, I open up Twitter, I feel so behind. And so I found out that you are a big skier and obviously climbing mountains. I would not call myself a big skier, but I do like to ski. And so my question for you, Jason, is what's your favorite kind of run to go down? If you say double black, we probably can never ski together, but what's your sweet spot? Give me this kind of run any time of the day.

Jason: [17:08] So I have two daughters, 15 and 13, and I've got to say they're incredible skiers. First run of the day will be a double black run. And no, I love steep and deep extreme skiing. I used to live for it.

DR: [17:19] Yeah, no. Well, that is not me. The only thing I do double of is double blue.

Jason: [17:24] I thought you were going to say a double shot.

DR: [17:27] The double blue is like the Goldilocks of runs for me. It's harder than just a normal blue, but the bumps and stuff, I'm like, my knees can't take it anymore.

Jason: [17:36] I can't do the eight hours anymore. So I'll generally, four to five hours and then I'll be in the bar at the end of the day.

DR: [17:41] Yeah, well, I'll meet you in the bar, that's for sure. We'll do some awesome après skiing. But I learned so much today about what Microsoft is doing and what you're seeing in telco and how people are starting to adopt AI. It's a fantastic topic that we'll probably be talking about for many, many years. So Jason, thanks so much for coming onto the podcast. Had a great conversation.

Jason: [18:00] I really enjoyed it. It's good to see you again.

DR: [18:01] Yeah. Awesome. Thanks. Stick around because we're ending each podcast with a Telco in 20 takeaway. I have 20 seconds to tell you something you need to know.

[18:13] Jason mentioned that the telcos that are serious about adopting AI are looking at their organizations in a whole new way, and they should be. You can't just apply AI to the way humans work today and expect game-changing results. While you might see some small improvements, in order to achieve the big results with AI that everyone is talking about, you have to do what I'm doing, redesign your organization with an AI-first mindset.

[18:38] Let me explain. Organizational design has its roots in the early 20th century evolving alongside the industrial revolution and the rise of large corporations. One of the first and most influential thinkers in this field was Frederick Taylor, who focused on improving economic efficiency and labor productivity through systematic studies of workflows and processes. He emphasized the importance of structure and management to achieve efficiency and goals. These theories were among the first to systematically consider the design of organizations and the role of managers in directing people and processes.

[19:11] But with the advent of truly useful and applicable AI, everything has changed. We need to completely rethink the way we design organizations. Instead of designing them around the workers, the humans, we now need to design them around the new systems that will be doing the work—the AI bots. We, of course, will still need the people to write prompts, tune the AI, and handle all the things that AI doesn't do well. But doing those things requires a different skill set than what companies have today. It's time to embrace the radical shift in how organizations will be designed and how work will get done.

[19:44] If you need a strategic HR thinker to help you reorg your telco, unfortunately, Jim Abolt's retired, so I'm your girl. Connect with me on LinkedIn and X at TelcoDR and we can talk about your awesome AI future. In the meantime, 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 amazing and well-read email newsletter on telcoDR.com and be sure to check out our killer YouTube channel. Later Nerds.