J.P.Morgan

MWC25: The Barcelona Express - The future of ubiquitous connectivity, AI and everything in between | J.P. Morgan | March 17, 2025

Shawn:

Welcome and thank you for standing by. I would like to inform all participants that this conference call is being recorded and parts of this call may also be reproduced in J.P. Morgan research. Views and opinions expressed by any external speakers on this call are those of the speakers and not of J.P. Morgan. If you have any objections, you may disconnect at this time.

I would now like to turn the call over to Samik Chatterjee.

Samik Chatterjee:

Yeah. Thank you, and welcome, everyone. I'm Samik Chatterjee, and I cover hardware and communication equipment at J.P. Morgan. I have the pleasure of hosting today the Barcelona Express - Future of Ubiquitous Connectivity, AI and Everything In Between with Amdocs. And with me are Anthony Goonetilleke, who's the CTO of the company, and Matt Smith from the Corporate Finance Group.

So what I'll do is we'll get into discussing the roadmap here for Amdocs, including how AI is increasingly proliferating their business. But before we do that, I'll hand it over briefly to Matt, Matt to go through the necessary disclosures before we get started. Matt, over to you. Thank you.

Matt Smith:

Yeah. Thanks, Samik. And on behalf of Anthony and myself, thanks very much for hosting us today. We really appreciate it. We will be making probably some forward-looking statements today, which are subject to the risks and uncertainties outlined in the SEC filings for Amdocs. And we may also reference some financial measures that are non-GAAP in nature, and you can find the reconciliation of those measures to the most comparable GAAP provisions on the Amdocs IR website and also in our 6-K filings. Thank you.

Samik Chatterjee:

Great. Let's get started then. Anthony, thank you, thanks for the time and great to have with you-

Anthony Goonetilleke:

Hi, how are you? Good to be here.

Samik Chatterjee:

You're probably better off to start with maybe talking about what were the key highlights from a Mobile World Congress, and we saw a lot of press releases, a lot of different discussions emerged from MWC, but maybe share what are the key highlights for you from both the industry perspective, what you found new and interesting, and when you think about what Amdocs showcased at MWC, what was that primarily sort of you would highlight to investors here? Thank you.

Anthony Goonetilleke:

Yeah. So I think MWC obviously is kind of a penultimate event for our industry from many aspects, and this year, it looks like it was back and forth. The stats I saw was over 110,000 plus people that attended, which is a very high number compared to a couple of the previous years. And so I think there was a lot of people attending. But more than that, I think there was a good buzz around the entire forum and the entire conference. I think people were looking at how do I grow revenue? What are my new opportunities? What are the new frontiers coming across? And obviously generative AI was at the front of everything. I don't think there was any hole or booth or stall or whatever you want to call it that someone wasn't talking about generative AI. So there was always important to sift between what's real and what's not real. But definitely, that was a very core theme.

Interestingly, there was also a manifestation of generative AI intersecting with physical devices. So there were a lot of physical devices this year that were incorporating generative AI, starting from phones to robots to dogs to cars to autonomous drones and all sorts of things. So it was interesting to see technology being kind of outputted in terms of different physical devices that are going to be practical in the next couple of years.

And of course, the intersection with all of this is the connectivity. So now with technologies such as eSIM where you can embed connectivity, almost everything had embedded connectivity. So you incorporate generative AI, embedded connectivity, the potential of next generation networks, all of a sudden all of these physical devices. I think I saw several different autonomous cars, autonomous drones, things like that where people were demonstrating all of these capabilities coming together. I think that was the really interesting thing.

Of course, there was a lot of discussion around the next generation of connectivity. If the last couple of years, people were kind of holding back on 5G standalone rollout and saying, "Hey, I'm pulling back on my CapEx spend." This year, people were like, okay, we've got the majority kind of launched out. I know there's been a slow rollout, but we feel like we're getting to a good capacity, a good level of coverage in terms of standalone network speaking to customers like Elisa that have 5G advance or 5.5 or whatever you want to call it deployed across. So now they're thinking of what are the next use cases we can deploy around a standalone network, which really, I know there's been a lot of discussion over the last, I don't know, five, six years around 5G, but you really can't get the true value of it until you have standalone deployed so you can deploy all of these use cases.

So that was really interesting and exciting to see people start to think about it. Also, our customers are starting to think about all these different revenue opportunities coming from added value resell. So there is generative AI capabilities to their B2B customers, to their enterprise customers, and also how to reinvigorate their enterprise customer base. This was a big theme I think that we saw come up everywhere.

Because I think if you think broadly on the consumer side, it's very digitized. I mean, today you can order a phone and you can essentially have it delivered to your doorstep fully activated by tomorrow morning.

It's not the same experience on the enterprise side. I know ordering the US, for example, takes you 35, 40 days to be from order to collections, so it's a very long period of time. And so there was a lot of focus on how do you accelerate enterprise growth, how do you focus on it? Because I think our customers look at that as a good growth market for them for the future.

So all of these things combined, I think it was a pretty exciting MWC. I know I come out of these things sometimes and you're so drained and you're exhausted. But this one, I think it was pretty exciting. I think a lot of action items, a lot of follow-ups with customers. Obviously, GTC is this week. I'm flying out for some keynotes out there to San Francisco tomorrow. And I think, so a lot of that dovetails in for Mobile World Congress to GTC with Nvidia to all of the other stuff that I've spoken about.

Samik Chatterjee:

Yeah. Well, helpful. And maybe before we dive into AI overall as a theme for MWC, what are the nature of conversations around 6G already when you talk about enterprise customers, and so there's enterprise opportunity that your customers want to leverage. Is that discussion about the next standards around 5G, or is that discussion more of a 6G discussion already? Because we often get this question from investors, when should we expect 6G, what changes with it? But is the conversation more around 5G, the next set of standards and enterprise opportunity, or is this more getting sort of pushed out to 6G?

Anthony Goonetilleke:

Yeah. So I think it's a spectrum, right? So if you talk about NSA 5G or non-standalone 5G, which was really the beginning of it, which is essentially just a radio is not the core network. Then you go to 5G standalone, which is then phase two of 5G. Now we have customers deploying 5.5 or 5G advanced, and in parallel rolling out broader coverage for standalone 5G at the same time. And then that kind of evolves into 6G.

So I think customers look at it as a continuum versus I think going from 4G to 5G was this massive big step leap. So I think we have these several steps down the line. And also if you think about the virtualization of the RAN in a similar way. So if you think of O-RAN as the big investment, the big kind of north star, there's many steps in between, right? It's splitting CU/DU, it's virtualization, it's moving to the cloud. So there's a lot of work, there's a lot of investment being done on that part before you even get to O-RAN.

So I think that's kind of how our customers are looking at it as a continuum and going, okay, what is my next ROI? What are the next things I can launch? Where do I now put my dollars in order to back that investment? So along with that are also a lot of opportunities around monetization.

So almost every conversation we had with our customer is like, "Hey, I need to monetize this differently. I need to think of bifurcated billing. I need to think about how you charge differently." Just being very creative on how you monetize different services, whether they be network exposure, because they want to get the maximum out of what they've already invested in, if that kind of makes sense.

Samik Chatterjee:

Yeah, yeah, okay. So maybe let's pivot to AI, which is understandably the biggest theme here. I think last year as well following MWC you highlighted that customers were already starting to ask about using generative AI and how businesses ask lower costs. How has that conversation changed since last year? What is sort of the discernible difference you're noticing in relation to the tone of the conversations with some of your customers here?

Anthony Goonetilleke:

Yeah. I think it's a hype curve, right? So obviously we are, I would say, some are at the peak of it, some are at the trough of disillusionment, some are... And not because there's any issue with the technology, but the thought process is, okay, so I have this, let's call it, I have this thing that can do these amazing things. How do I now incorporate it into my business?

So I'll give you a very simple example. So we have this platform called amAIz, which we've been showing customers, running pilots with customers, and the results are phenomenal. I have not seen results like this for... Actually, I can't remember the last time I saw results like this, especially in the care space. But now customers are like, well, how do I incorporate that into my customer journey? What is the impact on that on my people, right?

Because now suddenly you have this organization that had 10,000, 20,000 call center folks doing stuff, and they're going to be completely disrupted, displaced, the way they work needs to be changed. How you deal with a customer, the customer journey needs to be completely reimagined.

So I feel like there are a couple of different buckets. There are people that have done this business and worked in this business for a very long time, and they're kind of standing there scratching their heads going, "Huh? Okay, how do I incorporate this?" And then there are some people that are going to be the fast movers that are going to be, "Look, I'm going to reimagine my care experience." If you look at some of the results that we were showing at Mobile World Congress, we are getting... So there are three key parameters, for example, when it comes to care and call centers that people always monitor. It's average handling time, it's first call resolution, and it's NPS or Net Promoter Score, like how well did I do?

And generally, it's very, very difficult to have an improvement in all three. Why? Because if you're hanging up a interaction that I have with you very quickly, I'm usually going to call back because you didn't resolve my problem. Or if you hang it up very quickly, maybe I'm not happy with the result. Or if you don't hang it up and you stay long and you make me happy, then the average handling time goes up. So it's very, very difficult. It's like these levers, pick your poison, which ones are you going to tackle? We see with generative AI and kind of what we have that you can have a positive impact on all three. And we're not talking single digits. We're talking mid-double digits, like 30, 40% impact.

So with results like this, it's very difficult to argue, but then you need to go, okay, now, how do I incorporate this into my business? What is that change management or change management aspect of my business where I have to...? You almost, you hear this discussion if you listen to the Sam Altman's of the world, you hear this discussion, you almost need to reimagine what your new organization looks like. What does my org chart look like in order to incorporate these? It's not just a, here is a new tool, here is a new piece of software that does this. People that do that will have impacts on the fringes, but people that reimagine their new organization and say, "Hey, I have this completely new paradigm I can work with. What does my organization look like to handle that?" And I think those are the customers that you are going to see in the next 12 to 18 months jump ahead very fast.

Samik Chatterjee:

So maybe I'll take it a bit in a different direction in terms of just looking at the key parts of your AI strategy that you've outlined earlier. In terms of delivering that experience to your customers, I think the pillars you've talked about in the past are gen AI products, expanded data and AI services, and then POCs around the co-pilots that you're doing. Maybe if you can dive into the three, sort of the opportunities individually in those three pillars that you see.

Anthony Goonetilleke:

Yeah, sure. The first one to us is like table stakes, right? So every time we roll out a product now, we do four major releases a year, kind of our spring, summer, autumn, winter releases, all of them come with added co-pilot capabilities, GenAI capabilities in our product. So if you are buying our CatalogONE, every time you get take an upgrade, you're going to get more features, more functionality. You'll wake up one morning, and you'll have more stuff added into it.

So this is, for us, it's like we're disrupting ourselves, it's table stakes. If I try to draw a picture of what that looks like, if you've ever used our enterprise catalog or seen our enterprise catalog, we were so happy that it's drag and drop, and you can create offers and bundles, and see what it looks like, and test it and stuff like that.

Today with generative AI, I don't even have to know the system. I can just say, "Hey, Samik, let me create a new 5G offer in the Dallas metro area that's targeting millennials, and put this in this price point, and tell me what that looks like." And it'll automatically orchestrate and create four offers that you can then go and pick the one you like and tweak it. Whereas before, you'd have to pay...

Anthony Goonetilleke:

... right? Whereas before, you'd have to painstakingly create these offers and add the bundles and stuff like that. So this is where kind of that you're taking a two-week process and reducing it to five minutes, right? So this is where those capabilities for us, it's disrupting ourselves, it's table stakes. All our customers are going to benefit it the more they take Amdocs' products. Obviously, from catalog to ordering, to billing, to charging, they're going to enjoy these benefits. That's pillar number one. Pillar number two, or what we call type B, is really around the agents and the agentic behavior of we have these personas that are super, super trained in the telco-specific taxonomy and domains and understands the capabilities. So we have care agents, we have sales agents, marketing agents that you can incorporate into your system. Obviously, we'd love you to use an Amdocs system, but it doesn't have to be on top of an Amdocs system. So this is where we disrupt everyone else in the stream as well and this is where we have like over a dozen pilots running in different levels in production.

I think this space is very, very interesting because, number one, it's changing. If we had this conversation six months ago, it's very different to what it looks like today. You have things like MCP, you have things like where agents are talking to each other on a super agent concept. So I think this will continue to evolve and I think the biggest challenge for us, for our customers is that this space is evolving so fast, but it's definitely adding value. And how do you take something that's evolving so fast and incorporate it, back again to my initial comment about the org change management? So technology can evolve fast, people less faster, right? Because you need to train them. You can't keep changing organizations every month just because the technologies advance. So this is like a very delicate balancing act. But definitely, the results, like the stats I gave you are amazing.

The third part is kind of, it's interesting, because in a way it's kind of the dirty work, but it's also where I think there is a good opportunity in terms of a lot of services. A lot of the, I think Shuki, our CEO, calls it the plumbing, right? A lot of the infrastructure connecting these data lakes, connecting these data sources, all this integration, these interfaces, all of this needs to come together from lots of disparate systems, from mainframe, COBOL, all the way to full stack, brand new cloud native architecture. But they all need to feed essentially these generative AI capabilities. So there's a lot of services, a lot of integration, a lot of work, and so we can't ignore that. You can't just think of kind of the really cool, nice, sexy stuff on the top and ignore all of this kind of hard work. Where I think there's also an opportunity, obviously for Amdocs, to kind of do some work there.

So these are really the three pillars, and they go in parallel. It's not like one is better than the other, because of course we need to invest in our products and keep going forward. Of course, the agentic behavior is something that everyone is focusing on, which will revolutionize the industry. And the services are needed by everyone to fuel all of these capabilities.

Samik Chatterjee:

But when you think about timing, is there a different timing in terms of adoption by customers of these? Is there primarily like the Copilots are the ones that you would expect to get adopted first, and then followed by the other two? How you thinking about near-term versus more longer term adoption?

Anthony Goonetilleke:

Yeah, the adoption of part one, like you take our product today, you're going to get it and you're going to start using it. So the adoption is almost like symbiotic, it just happens, right? The second one is a little bit more difficult because here you need to rethink your organization and change your organization. The third one, we are already starting to do and we're doing it because you have to normalize the data. You have to understand your data. You have to understand what the logical model is that you're feeding. At the end of the day, the old adage of garbage in, garbage out is very relevant. So you need to ensure that you are not training your models, you're not using data for rag or whatever other purposes that are not accurate.

Just as a simple example, obviously we're disrupting ourselves internally as well, and we had a, someone had built an internal corporate system around HR, and I was playing around with it. And I was looking at vacation policy, and it was giving me different results. And of course, why was it giving me different results? It was because the information that was fed in that was used was just outdated. Right? So this whole layer of what you put in becomes so critical. Your policies, your procedures that may be set on a shelf becomes much more relevant to keep alive and current in order to adhere to the business, especially when you're getting generative AI to make the decisions.

Samik Chatterjee:

Yeah. Maybe on a different approach. I mean, you talked about the improvement on all the scores that a customer cares about, right? Like NPS score improvement, there's also better cost management or lower cost. But when you're talking to the customers, which innings are we in, in relation to AI or convincing them of the value of AI? Is it still in the phase where you have to go out to a customer convince them of the cost saving or the improved performance? Or is there a different level of discussion now that the cost benefits or the performance benefits are well understood, now it's more of architecture decision of what needs to be done to implement it? Or even what are you hearing if a customer says, "Okay, I have certain hurdles that I need to get over before I can implement this." What are the typical hurdles that you're running into?

Anthony Goonetilleke:

Yeah. Look, I would say that almost every one of our customers has some type of thought process or strategy around generative AI. What they are struggling with is they are being bombarded by 50 different companies telling them 50 different things. And saying, "Hey, use my product," or use this or use that. Plus, they have a whole bunch of companies that are essentially going to be displaced because of generative AI. So if you were a BPO company or if you're a call center company that charge per seat or something like that, I mean, your business model is completely disrupted, right?

So I think there's all of these forces at play, which if you are a CEO or a CTO or a head of technology, you're listening to all of these different voices and you're trying to figure out, okay, what actually makes sense for my business? On the other hand, your head of HR is saying, okay, how is this going to impact our people? Are there unions involved? What does it mean? What does it mean for our offshore strategy? All sorts of things. Because remember, there are a whole bunch of companies that suddenly used to make money by you hiring more people. There's a vast amount of companies, if I increase my number of people from 500 to 5,000, I get more money, right? That world is going to be completely disrupted if that was your model. And so I think there are companies trying to grapple with it. There's a lot of smoke and mirrors out there. Today, it's very easy to just go to a large language model and put something in and spit it out.

I'll give you an example of this. So we showed our maze platform, and one of our customers who was building something internally, said, "Hey, we have something like this as well." And they were basically feeding entire bills into a large language model and getting results. And one of the results was, "Oh, by the way, if your bill is too high, go check other service providers." Right? I mean, it's not just about feeding something and hoping that someone else trained it and you get the right results. Because there's a cost aspect of it. Tokens are expensive, GPUs are expensive. CFOs, from the last time I checked, I know our CFO likes predictable models. If you're giving your entire organization this toolset and you have no idea how much they're going to use it, suddenly you get to the end of the month and you're like, "Oh, my gosh, what the heck happened here?" Right? You're making one of these infrastructure companies get very rich, but you have no predictability. So all of these things need to be taken into account.

So it's not just about rolling out generative AI, it's about thinking of your ROI, it's thinking about your organization, thinking about what that model looks like. So this is why I think there's a little bit of, I would say in the next six to 12 months, there's a little bit of shifting going to be happening where the folks that are doing the real stuff are really going to stand out. And the companies that are really trying to defend their territory because they were charging per seat, per head, per BPO, are really going to start to fade. And I think this is the kind of the intersection we're at now.

Samik Chatterjee:

Maybe on that front. When you think about customers leveraging Amdocs, and particularly sort of relative to this AI approach that they have, are they opting for a full-stack solution, going from the very sort of Copilots to eventually the AI product that leverages some of the expanded data and the AI services that you provide? Or are customers more saying, "Let me try one thing. I'll add to that stack from Amdocs later if I need to"? What is the approach customers are taking? Are they eager to sort of dive in and adopt the full stack or more disaggregated on that front?

Anthony Goonetilleke:

Yeah, so it's an interesting... We also had an interesting crossroads, I think, around retiring tech debt. So if you think of legacy systems and mainframe systems and the whole VMware saga and things like that, that's occurring, there's a lot of our customers are going, "Okay, what does my next generation digital stack look like?" Because generative AI works hand in hand with your digital experience, right? So these two go hand in hand. So customers are always reimagining.

Now, you need to be able to capture an order, you need to be able to handle it. Remember, our industry, I always give this example, if you are an e-commerce vendor, like, I don't know, pick one, let's say you're Zappos, you sell shoes. I come in and say, hey, I want this pair of shoes in this size. You go, "Fantastic, I have it in my inventory. I will ship it out to you." Transaction done, closed. If you're a customer of a telecommunications company, you come in and say, "Hey, I want this phone." "Okay, great." "I want to port in my number for my previous phone. I need an e-SIM provisioned. I also need to add it to my family plan. Oh, and I'm canceling my other phones, so I had a early termination fee for that, so that needs to be accounted for. And I need to be prorated because I'm buying a bundle now and I'm adding four phones." So all of this complexity when you're buying a phone needs to be completely seamless, right? So this doesn't go away.

So our customers need to have a full stack in a way is almost deepened, because having 70 different legacy systems and legacy networks and things like that kind of prohibits you from doing things, right? So let's take generative AI and let's take our catalog one example. If you take our new catalog one product and you're using generative AI, not only can I give you the answer, I can orchestrate it, create it, and provision it for you. If you don't have the full stack, maybe you can get the answer. Someone needs to then copy it, print it out, maybe look at a system, okay, so it's this, drag it, drop it, then hope it gets provisioned. So you get all these disparate points around the stack. So I think when our customers look at it, they understand it, right? Our customers, we still have a lot of digital transformations going on. I don't think that's going away. I think generative AI is just going to augment it and get people to rethink what that picture looks like or what that future looks like, I would say. The customer journeys are definitely going to change.

Samik Chatterjee:

Maybe before we go down any further in terms of the platforms, et cetera. One, as a reminder to the audience, you do have the option of submitting a question through the platform you're logged into. So in case you want to avail of that feature, please do. And we can take your questions over chat.

Maybe, Anthony, before we get into the more specific aspects. One thing is your customers traditionally, these telco customers are not really known to be the best in terms of being on the cusp of the technology curve or leading edge. And so what are the typical hurdles that you are seeing? Like when you talk about being prepared for deploying AI, obviously the customer needs to have their legacy systems or their data in a certain very concrete format where it can be easily digestible by a model. How much of a hurdle is that or how much of a sort of overall staggered deployment of AI do you expect because of that primarily? Is the ground perception that most investors have, that these telco customers are not necessarily the leading edge in terms of how their data is organized? Is that true and is that a hurdle that you need to overcome over time as well?

Anthony Goonetilleke:

Look, I think it's maybe a little less of the organization of data, but more about how they utilize it. So obviously, being in the telco space, there is a lot of PII, privacy issues that you need to take into consideration, right? Where you put your data, how you use it. And I think telco as an industry are very, very conscious of this and aware of this and understand there's regulations around it, because at the end of the day, this is very sensitive, very private information. So I would say there's that aspect of it.

The second aspect of it, I would say is they think about... When you think about how fast you move and what you do and things like that, it's yes or no. If you think about cloud for example, telcos were probably fairly slow to embrace and move to the cloud historically, if you go back and look at it. And people are like, "Oh, they were slow industry to move to the cloud." Now, we realize there was a reason they were a slow industry to move to the cloud, because telcos, at the end of the day, if WhatsApp is...

Anthony Goonetilleke:

... the cloud because telcos, at the end of the day, if WhatsApp is down or Gmail is down or Snapchat is down, people are like, "Eh, such a hassle. It's down." If AT&T or Verizon or T-Mobile is down, it's like, oh my gosh. It's like front page news. It makes the Wall Street Journal. So the dependability of the population on telecommunication service is so high. It gets confused with the slowness of an industry moving. So if you look at 5G, and if you look at the rollout of 5G, and if you look at how fast these standards were adopted, and now you see, I look at countries like Singapore, South Korea, we spoke about Alisa before, how fast some of these countries are moving. And of course in the US standalone, it's a huge embrace of technology and they're in control of it.

So I think same with generative AI, I think you get a bifurcation of not just saying, "I'm going to take all my data and put it into some unknown large language model that's out there somewhere and trust it" versus "What do I put in? What do I control? What do I keep? How do I wall fence it? How do I..." All these types of questions and they're the right questions to be asked, but I have to tell you, everyone is head down in this. I don't think there's any one of our customers that are not thinking about it, that are not doing... I mean just with the POCs that we have running with our customers, almost everyone I would say is in some shape or fashion thinking about how I incorporate it.

I think the key again is going to be... The tipping point is going to be, look, it's hard to argue what benefits it'll bring from a customer care perspective and a cost reduction perspective. You have no bigger lever on the table to reduce cost in an organization that is very people heavy right now than using generative AI and it's material, the cost reduction aspect.

So if your CFOs or CEO is going, "Hey, I need to take cost out of the business." This is the tool that you have. So I think that's going to drive some of the faster adoption on it versus you take something like cloud, you needed all of these requirements, plus cost-wise at the end of the day may or may not have been cheaper or not cheaper. It's six of one, half a dozen of the other. Of course, you've got a lot more benefits from moving to the cloud. So I think the adoption of this will start to increase gradually. But then the other thing is also really, and I keep mentioning this, is really the heavy people side, the unions, the customer care, the outsource call centers, all of these things need to be rethought.

Samik Chatterjee:

That's clear. Maybe moving to your platform, Amaze, what's been the feedback since you rolled it out to customers? And when you look at from Amdocs' perspective, what's the roadmap for that technology or enhancements on that platform over the next few years? How are you thinking about it?

Anthony Goonetilleke:

Yeah, so Amaze is basically a framework for us that we kind of built all of this telco taxonomy, trained it. We added these governance modules to ensure data security, PII information is secure, it's safe. We added kind of studio capabilities working with some of our web scale partners so that you can start to build the use cases on top. So all of these different components are boarding together versus you having to now go and work with 36 different open source companies and try to stitch it together. So this is really what the framework is. And it's purely, I would say highly, highly focused on the telecommunications cable sector. And so from that perspective, I think our customers like it because it's very, very focused on the industry and solving the problems. You take something like early termination of a device. We know for example, because we've been in the industry and because we have the full stack systems, I can tell you that there are over 200 different places in different databases that impact early termination and just using a large language model, it doesn't know that.

So this is kind of the knowledge and the expertise I would say we bring to the table and also making it cloud native. So of course we have our partnership with NVIDIA and our partnership with Microsoft, of course working with AWS as well and Google. But working with the web scale guys hand in hand focus on the telecommunications industry I think has been a benefit for us. This is why I think we see all of these pilots taking up and customers going, "Oh, okay, so I don't need to build this component. I can just focus on the functionality versus building all this infrastructure."

And outside the US we also see an interesting bifurcation of people looking at this sovereign data center concept of potentially maybe having generative AI reusing their data centers to maybe build out GenAI farms and GenAI capabilities to even sell to their enterprise users. So we have one customer that's using our Amaze framework that is basically using a framework on top of, they're right now in the process of building out a GenAI farm to sell to their enterprise customers because they want to have data sovereignty in country. They want to have privacy in country, and they believe they're in the best position to provide it. So it'll be interesting the next several years, this will span out, but I think it'll be a hybrid kind of topology that will evolve.

Samik Chatterjee:

Since we are on that topic, are customers showing any concern about taking their data to the AI or concerns around that proprietary data that they have, but making it accessible to the AI models? Have you seen any concerns from your customers?

Anthony Goonetilleke:

Yeah, it's very high, right? I would say the telco industry probably maybe outside of HIPAA stuff is probably one of the highest in terms of the concerns. But they're addressing the concerns by, they're not using just standard public models. Well, not the serious guys anyway. So they're either using instances dedicated for them, they're using open source LLMs that are very specific for them. It's a big concern, I would say. And the Amaze framework also allows you to split transactions across multiple large language models. So for example, you take something like OpenAI, which is probably still the most amazing tool to dissect natural language coming in from a user and understanding what the breakdown of the question is, right? It is probably still the leading one.

But on the other hand, once you've broken it down and segmented it and the execution and the understanding of the different elements, you don't have to necessarily only use OpenAI. So we use several different large language models. For one transaction, we may use three different models and send different components to different models, and that way we can maintain privacy, not transfer PII information, where else we may use a standard model just to understand what is the question that Matt is asking, for example.

Samik Chatterjee:

Maybe let's talk a bit about use cases. You've talked about this a bit already, but if you had to give specific examples of what use cases the customers started with looking at in terms of leveraging AI for and now two, three years down the line, what are these sort of additional incremental use cases that they've been able to think of that where they can leverage AI?

Anthony Goonetilleke:

Yeah, look, I think care is the big one. The agentic care persona, I think is the poster child for generative AI. If someone knew nothing about generative AI, I had no clue, and they said, "Tell me what real practical example that this can do," I would show them an agentic persona of customer care, because this is revolutionary, and even I think the word revolutionary is underserving to the results that it provides. I think those use cases are maturing. So take telco for example, probably around 30 to 40, maybe even higher percentage of calls coming in are people going, "Hey, there's something wrong with my bill. The cost is higher. I have this extra charge. Why is this here? Why did my disk..." Some query about the bill. Right? And when you can materially impact this, when you can potentially say, "Hey, I can take your average 14-minute call and reduce it to 45 seconds."

It's a complete game changer and we see this improving. So last year compared to this year, this year, we're getting 97, 98% accuracy. I would challenge anyone that humans are less accurate than 97, 98% where people are not following policies, procedures, not understanding the company discounts what they're allowed to do. So I would say if you apples to apples comparison, humans are much less accurate than that. So I think the level of accuracy, the level of trust, the level of... Another thing we haven't spoken about is around latency. A lot of the interactions, a lot of what we do needs to be in near real time. There's not a lot of industries that need near real time access to capabilities.

So when we're building this stuff out, these are some of the principles we are keeping in mind. So these are some of the things that continue to evolve, like accuracy is going higher and higher compared to last year, latency is going lower and lower. And this is why, by the way, our framework also broke down those large language models and small language models so we can get lower latency and better accuracy because obviously some small language models are very trained for very, very specific items.

I would say that and around the governance. So what are you allowed to do? What is it can you answer? What can you not answer? And we all know the examples that have been out in the media. Every time someone provides any type of generative AI capability, people are always trying to ask it stupid questions and get funny answers, and then it'll be posted on the news the next day. So putting guardrails to all of this stuff and not allowing someone to come and get dating advice from your care customer agent in the telco space, right? It's funny, but it's serious stuff. And these are the things I think that are really advancing and really, I would say a game changer. And we'll have a chat in a year from now, and I think the advancements will be amazing. I think the next big thing to look out is the interconnectivity between these agents. So today it's very early days, but I think in the next six months you'll see this interconnectivity and the orchestration of this just advancing in leaps and bounds.

Samik Chatterjee:

But maybe let's move to talk about your customers and their appetite to invest in this capability themselves. If the customer was to think that this is something I can build in-house, one, how feasible is it? Secondly, how many customers are even showing the willingness to do that?

Anthony Goonetilleke:

Yeah, I think everyone has some team in-house that's playing around with it, but I think at the end of the day, we are going to move from experimentation of going, "Hey, you are paying for all these GPUs. Now show me some results." So we are going to have to move very, very fast. And one of our customers had built a bunch of staff and they said, "Hey, come in." And they called it the Pepsi Challenge. I'm like, "What do you mean Pepsi Challenge?" He goes, "Pepsi-Coke Challenge. Let's run it side by side and look at the results." And they ran our side-by-side with it and it was a no brainer. Why? We are investing millions of dollars in R&D to kind of mature this bill, the telco taxonomy, whereas they're just taking a bill and uploading it to a large language model and hoping the results will be okay.

So I think we are going down to this path. I think the investment will be there. What I am looking for is this kind of tipping point. There's going to come a tipping point. So right now, white papers, people are talking about it, people are talking about agentic. People know that there is benefits for it, but all of a sudden, you're going to have 1, 2, 3 of our customers that are going to do it, that are going to execute it, that are going to come out publicly and say, "We were able to cut our cost structure by 35%," and then it's just going to be like this flood. So getting ready and being there at the right time, having these pilots in place, being able to show that they can do it, because the moment one does it, the next question, the CEO is going to call up and going, "Hey, I heard X did this. Why can't we do that?" And then we want to be in the position where I'd say, "Actually we're trying this out. We think we can roll it out in the next three months."

Samik Chatterjee:

So on that part, let's say you do have this scenario where more and more customers are coming as soon as you hit that threshold, but how much of the existing stack that can you leverage one customer to another? What is the level of customization? How has that time to deployment with a new customer changed for you? I mean, you've obviously gone through a learning curve yourself, I'm assuming that. Have you been able to materially impact the time to deployment for a specific customer?

Anthony Goonetilleke:

Yeah. So the Amaze framework, we can deploy it today, six to eight weeks to production, not just testing, to production. The long tail is not the platform or the framework. The long tail is the integration with the data sources. So when I say six to eight weeks, we can deploy it to production, this is with maybe about 10 data sources. Now obviously our customers have dozens and dozens of these. This is really where the long tail comes in terms of how quickly you can move and deploy it and where a lot of the work. Back to you, your comment you made before about the plumbing and the services aspect of it. There's a lot of this integration that needs to be done to connect it. In terms of the framework, though-

Anthony Goonetilleke:

... To kind of connect it. In terms of the framework though, this is where we're investing in the product. So it's like one framework, it's one platform. We have it running on a couple of the clouds. We also have one version running on-prem. We work very closely with NVIDIA. We were one of the first companies to work with their NIMs infrastructure, for example. We'll share more about it at GTC this week. I know last year, even Jensen mentioned it, hopefully we'll see how it goes this week. But I think keeping very close to the likes of NVIDIA and Microsoft while evolving our platform is less about customization. The customization and the work comes about integrating all your legacy data sources and normalizing that data and connecting it in. That's where a lot of the work is.

Samik Chatterjee:

Maybe let's talk a bit about cable customers. I mean, telco customers are your core, I would assume. But when you think about cable customers, how should we think about attraction and more importantly, are there a different set of use cases to think of?

Anthony Goonetilleke:

Yeah, not really. I mean, when we talk about our communication service providers, we include telco, wireless, wireline. We don't necessarily look at it very differently. The customer care, by the way, there are three, I would say primary categories of customer care. There's an issue with my bill, there's an issue with my service, or I want this new service, like a sale. Those are the three primary and they span across cable and wireless and all our customers. I would say the interesting opportunity is also around network and service assurance. So there are some very interesting capabilities. So we have a service assurance platform that we acquired, I'm trying to remember when we acquired it, but maybe a year and a half, two years ago now around service level management, performance management, outage management and things like that. So all of those, rather than the legacy way was to have dashboard and get alerts and things like that.

Now you go into a predictive model of going, "Hey, here are my seven leading indicators. This is going to lead to a bottleneck within the next half an hour." So those are some of the interesting use cases on the cable side, I would say that we are looking at right now and design partners together with our customers, looking at how we can help them. Also, around some of the, we're working with a couple of the satellite companies for example, because connectivity today, I think of connectivity, I call it ubiquitous connectivity. Nobody cares where you're getting your connectivity from.

My daughter had a school trip down, they went down the Amazon River last week and she was telling me how she got connectivity and she was getting D2C connectivity on her phone. It wasn't that they had cell phone towers, so she was able to send a text and stuff. Whereas my eldest daughter who went five years ago, I bought her a satellite phone to take on that same trip. Today I didn't have to. So when we think of connectivity, again, it's a spectrum and all the way from D2C, Leo, MEO, and Geo satellite plays to 5G-Advanced, 6G, 4G, LTE to fiber, to cable to DOCSIS, and you have this spectrum. At the end of the day, the consumer doesn't care. So how do you look at this network holistically? How do you problem solve? How do you flip between one mode to another mode?

Another thing we're doing, this is one we were working with NVIDIA on, is around digital twins of towers, on physical towers. So outages and things like that. If you have a digital model and if you have an asset registry that feeds into our OSS, can you model this from screws rusting to you've got a bird hit and something's gone completely wrong and knowing where it is? Because a lot of the problem is also I have these towers in the middle of nowhere sometimes and someone needs to go there, they need to figure out where it is, they need to know where power is, they need to know how to problem solve. So kind of these digital twin capabilities and the generative AI capabilities, not just spans from the use cases we know up the front, which is just kind of like the cool ones, like customer care, but all the way to some of the big back hole stuff, the spectrum optimization, the service assurance use cases, which shortens the time.

I would say another interesting one our customers are looking at is revenue recognition acceleration. So if you take an enterprise order today that takes 35 to 40 days, if I can shorten this to two weeks, to one week, you suddenly have a revenue acceleration opportunity that you can capitalize on. And definitely generative AI plays a big part here. So if I was an enterprise sales rep that went and sat at a customer and said, "Hey, what do you want? You want this? Okay, let me go back, let me see if I can..." And this goes back and forwards for a couple of weeks. Now using generative AI, I can check my inventory, I can check my sources, I can build your proposal while I'm sitting in front of you, at least do 75% of it. So this helps you accelerate revenue. So this is another, I would say a good opportunity for them.

Samik Chatterjee:

Before, I know we are getting close to the end of the hour, so maybe let's pivot a bit towards talking about network automation in relation to RAN platforms and size, the timing of the opportunity. And I know you also had an announcement with your partnership with Google to enable distributed 5G on cloud-native technology. So maybe let's go through those two things. Firstly, how do you think about the automation and its application on a RAN platform and the size of the opportunity there? And then we can touch on your partnership with Google.

Anthony Goonetilleke:

Yeah, so I think the virtualization of the network is in full steam and it's going all the way down to RAN. Now you can argue is O-RAN going to happen? Is it not going to happen? I'm not even going to get into that philosophical debate. It may happen, it may not happen. We'll see where the NEPs go with that. Obviously they each have their own view on it from a slightly different perspective. But the virtualization of even moving a single vendor to be virtualized and splitting CU/DU and moving to the cloud, I mean, that's real and that gives you a lot of benefits that you can optimize spectrum, that you can get data insights. So I think in the next several years there will be a lot of spending in that space to virtualize it and that will potentially be the first step to maybe a future of O-RAN in the future.

But I don't think you need to either believe in O-RAN or not believe in O-RAN to understand that virtualization, I mean, we have customers where virtualization is happening today. We did a project with, I think this is pretty public, we did a project with Dish and AWS where we helped them virtualize their RAN. You can do it today, you can run it today. The technology's there, it's solid. So I think that would be a multi-year opportunity and I think it'll be very interesting.

I think going to the Google announcement, back again to service assurance and everything you can do from SLA management to performance management to metric calculation and doing this in a closed loop and using this data to predict potential future outages, potential future spectrum allocation, this is the thing that we're partnering with Google and looking at. We did a demonstration actually in the Google booth. We were showing some of these capabilities in Mobile World. This is something that's pretty exciting, because now suddenly if you think of an outage and you think of all of the lost revenue and the potential there, the fact that you can get ahead of it is pretty exciting.

Samik Chatterjee:

Maybe I'll wrap up with one question, which is essentially what investors always look for, which is trying to quantify what it means for growth and margins. When you put all these opportunities together, I know for 2025, fiscal '25, you're guiding to about 2.5% revenue growth, 21% EBIT margins, 21% plus EBIT margins. What does this opportunity mean and is it more impactful on the way you think about Amdocs as revenue growth in the future? Or is it more impactful on what margins Amdocs is able to own in the future?

Anthony Goonetilleke:

Yeah, maybe I'll start and Matt can jump in here. Look, I think if you look at what we've done in our margin profile, definitely we're baking in, we're disrupting ourselves internally. Now, it's obviously not a one-to-one correlation, because there's customer contracts, the customers expect savings and things like that, but you can do kind of some back of the napkin here and you can understand that it's split between both. Giving our customer savings and obviously on the margin profile.

This will continue. I think this will continue in terms of we have to be more efficient, we have to work better, but some of this is also delivering more capabilities to our customers, which is very important. So our customer may come to us and say, "Hey, I paid you $10 and you delivered 10 widgets to me. I want to pay you $10, but you are now going to deliver 15 widgets to me." So some of this comes in that form as well, where we deliver more epics, more PIs. If you think of more kind of scrum-based work where we can deliver more value to our customers, I would say. Matt, I don't know if you want to jump in here in terms of some of the other stuff.

Matt Smith:

Yeah, I mean, in terms of the market opportunity, we've talked about a service addressable market of going on $60 billion this year. We're a 4.5 billion plus company. So within that market opportunity, there's a huge potential for growth across the five growth pillars that we have. Digital, 5G, cloud, B2B, and so on, and generative AI is running through all of these as a theme. So I think we feel very confident in the market opportunity and the room to growth. We are seeing some acceleration in terms of the guidance for Q2, although the full year outlook is 2.7% pro forma revenue growth. In constant currency terms, we're already picking up a little bit in Q2 and we're expecting about 3% year over year if you take the midpoint of the guidance. So I think we feel good about the revenue growth outlook around those levels.

And in terms of the margin, we've talked about ongoing efficiency gains as being a core component of the margin expansion that we expect to deliver this year. We did have some phase out from some legacy activities, which is driving about 230 basis points of improvement this year. But another 60 or 70 points is coming from the things Anthony was just talking about, these internal efficiency objectives, the cumulative benefits of which are really beginning to come through in the business and driving that margin expansion this year. And we think margin expansion beyond this year is sustainable. I think it's probably a bit early to quantify it and to give numbers around it, but certainly as we continue to implement automation, traditional AI, machine learning, and then with the benefits of gen AI, we feel that sustainable margin expansion is something that we can deliver on.

Samik Chatterjee:

Thank you, Matt, Anthony, thank you for the time. We're up on time here, so I'll take the opportunity to wrap up the call here. But thank you again. This was very insightful. Thank you both and thank you to the audience for joining as well.

Matt Smith:

Thank you very much for hosting us.

Anthony Goonetilleke:

Thank you so much. Bye.

Samik Chatterjee:

You too. Bye.