05-Jun-2026

Cognizant Technology Solutions Corp.

(CTSH)

Cognizant AI Forum

CORPORATE PARTICIPANTS

Tyler J. Scott
Senior Vice President-Investor Relations, Cognizant Technology Solutions Corp.

Ravi Kumar Singisetti
Chief Executive Officer & Director, Cognizant Technology Solutions Corp.

Rohan Narayana Murty
Co-Founder & Chief Executive Officer, Workfabric, Inc.

Nitesh Ambastha
Executive Vice President and Group CIO-Firmwide AI, Digital Platforms & Investor Experience, LPL Financial

Jatin Pravinchandra Dalal
Chief Financial Officer, Cognizant Technology Solutions Corp.

Ravi Kiran Kuchibhotla
Chief Strategy Officer, Cognizant Technology Solutions Corp.

MANAGEMENT DISCUSSION SECTION

Tyler J. Scott
Senior Vice President-Investor Relations, Cognizant Technology Solutions Corp.

All right. All right. Thank you, all. Thank you, all. Good morning, everybody, here in the room in New York. Before I get started, go Knicks. Anybody? Yeah. All right. A special hello to everybody on the webcast as well. Thank you all for joining us and welcome to Cognizant's 2026 AI Forum.

My name is Tyler Scott. I'm the Senior Vice President of Investor Relations. And it is a pleasure to welcome you here today. We're excited to be joined in person by many of our industry analysts, advisors, clients, partners, and, of course, some of our shareholders and sell-side analysts. So, thank you so much for being here.

Today's AI Forum addresses a defining question facing businesses right now. How do you close the gap between the AI promise and the potential and the enterprise-wide impact? Last year we put our AI Builder strategy in motion and bridging this gap is exactly where Cognizant operates. This morning's keynote session is going to be webcast live with our CEO, Ravi Kumar; and our CFO, Jatin Dalal. And a replay will be available on our website after this call. For those of you joining in person, we encourage you to take advantage of the networking, demos and engagement opportunities in our sessions later today.

And now, of course, one of the most important updates. While we're not going to be providing financial updates, some of our comments on today's webcast and in our meetings today will be forward-looking. These statements are subject to the risks and uncertainties as described on the slide behind me. And, of course, our other filings with the SEC.

With that, let me welcome our CEO, Ravi Kumar to the stage.

Ravi Kumar Singisetti
Chief Executive Officer & Director, Cognizant Technology Solutions Corp.

Thank you, everyone, for making it today. The weather in New York has got better in the last few days, so very excited to host you all. We have a very diverse audience, sell-side analysts, industry analysts, some of our partners, investors. And most importantly, while we speak about our strategy – in fact, when we got to an Investor

Day last year, we spoke about expansive margins, being on the winner's circle and our EPS growth being greater than our revenue growth. So that mission and that aspiration continues to be with us.

Most importantly today, we have our clients who are going to talk about some of our early success, learnings, early wins, the partnership with Cognizant. LPL Financial, Citizens Bank, Signature Performance, Elevance Health, U.S. Bank, Lineage Logistics, Sysco Foods, TD Bank. So, we have eight clients talking about early learnings, early wins with AI.

Every company and every industry is set for a purpose. It's set to address their white space, a problem – a purposeful problem and companies are built on a mission around it. For two decades our industry, Cognizant built a mission around expertise to navigate enterprise transformation with the power of technology. Classical software manifested into systems and we globalized and scaled companies by applying those systems to business flows. And we built a craft.

The craft was about building the systems to scale enterprises to efficiently and effectively run. And we could do that better than our clients because of the repeatability of the rinse and repeat of what we do for multiple clients. And we embed that into software engineering for systems and products. And we kind of got into the adjacencies of managing the infrastructure, managing the process, managing – over a period of time infrastructure went virtual, so we managed to transition ourselves through this tech disruptions and kept that unique craft of helping clients go through that transformation.

We transitioned ourselves from being a system builder to a systems integrator over the last 25 to 30 years. Here we are now with this extraordinary AI technology, it's a platform shift, Internet scale data, neural networks, and an extraordinary compute all coming together. And it kind of blurs the line between systems and people to run business flows. And because it blurs the line, we have a much bigger sprawl beyond systems to operations.

Look at the economic data. This is a tailwind of what we do to the AI opportunity. $20 trillion of economic output by 2030, of which $6 trillion is going to be productivity. 14-plus trillion dollars, so one-third – two-third, $14 trillion is about new products, new services powered by AI. A lot of what you see today is productivity-led. And this is going to create some kind of a flywheel. There's $1 trillion of spend on the AI infrastructure, $6 trillion to $7 trillion is going to be spent in the next three to five years. And effectively that output is only going to come if the capability of AI, which is right out there and it's going higher and higher, it's only going to get better and better and the production value to enterprises is way below.

And that is the gap we are trying to address. $1 trillion already invested, another $6 trillion to $7 trillion to be invested by 2030. The capability is going up. The production value is right out here. In some ways, that's the gap we are going to address as a company.

The costs are ballooning with very little productivity and we're already seeing enterprises talking about it. The challenge – I mean, MIT spoke about it, McKinsey spoke about it, Bain spoke about it. Uber put it on the list – on their earnings, Microsoft spoke about it. I saw a report of Walmart actually speaking about how this is becoming a challenge. Sam Altman last week spoke about how he's already seeing budget issues related to AI in enterprises. So, there is a lot of reasons why the production value is way below the capability. If chapter 1 was about broad open based experimentation, chapter 2 would be about specifics, realism, costs, and bridging that gap.

And here we are building a new category of a company which is going to evolve from the past, some craft from the past, some new craft. We would now build this new category of a company with a new purpose to bridge that gap. And that is what I think the opportunity in front of us is.

All right. So, we did this study called New Work, New World in 2023. And then we then we did a revision of it in 2026, 18,000 tasks, 1,000 occupations from the O*NET database. And we actually projected that 10% exposure to AI – 93% of the jobs would be 10% exposure to AI; 25% exposure to AI, 69% of the jobs; and 50% exposure to AI, 30% of the jobs. And this specific exercise, which we did, we thought all of that is going to happen in 2030. Here we are in 2026. It's already there because the velocity at which this is going is significantly higher. So, we did exposure scores, friction scores and velocity.

And what is really happening is systems plus people which created flows and our universe was systems is now becoming systems plus people plus digital labor or agentic work. And the universe is completely now spread across this entire spectrum. And the lines between systems and people are getting blurred.

And I'm going to give a little bit of a back-of-the-envelope math on this. And that's where I got this $4.5 trillion, too. The total global economy is roughly around $120 trillion. Global 2000 companies have $52 trillion of revenue. $20 trillion is operations related spend, $1 trillion is spent on systems by the global enterprises, G2K, they spend their $1 trillion on IT services, system integration work, they spend $350 billion to $400 billion on software – enterprise software. And now the universe we are talking about is $1 trillion of IT services from systems. We'll talk about it, it has a new flow, and we have one-third of the $15 trillion to $20 trillion which is invested on operations of companies by the Global 2000, which is $20 trillion, $4.5 trillion is the labor exposed to agentification.

So, the universe we are talking about is $5 trillion to $6 trillion. So, it's a 6x opportunity to the opportunity we actually chased for the last 25 years. That $6 trillion of an opportunity is going to be labor in business operations of companies, agentification and rewiring the systems to do that.

And when we do so, you're going to create a flywheel because as the labor gets agentified, the human endeavor is going to look for new purposeful work, and that flywheel will create consumer spend and the $20 trillion of economic output will create more wages and more jobs. And we think those jobs of the future are going to be significantly higher than the jobs of the past. In fact, there's a World Economic Forum report which talks about how [ph] 125 (00:11:18) million jobs are going to be created by AI by 2030, while 70-odd million jobs are going to be knocked away from the past.

So, what are we trying to do? We are trying to find new value pools in this new universe of $6 trillion, and that unlock of the new value pools is what we're going to look for as the opportunity to bridge that gap and the opportunity to add value to actually bridge that gap from capability to production value for enterprises.

The pace of change has been absolutely high clock speed. There's been a sense of FOMO, fear mongering, and that has led to token consumption without linkages to ROI and without linkages to outcomes. So, one of the reasons for this gap between capability and production value is also because there has been relentless token consumption without linkage to outcomes.

Let's take the frontier model companies, roughly $100 billion of annual revenue as of today. They're probably going to do $1 trillion by 2030 – $500 billion to $1 trillion, if the scale-up of the infrastructure continues the way it is continuing. Now, let's take a more optimistic scenario, $1 trillion. That revenue of $1 trillion because of the effectiveness of how the token consumption should be, a part of it is actually going to be routed through system integrators or AI builders, as I call it, like Cognizant. Why? Because you need to create more efficient, more effective – remember, this is a contextual science, more predictable and better economics for the token consumption.

So, Cognizant is actually building a token and an agent harness, I call it, with model operability, interoperability. The ability to capture the patterns of work, the ability to wire the digital traces and actually create a compounding factor for the digital labor to be routed through us. We do this on a rinse and repeat for hundreds of clients. So, shouldn't we be doing better than our clients because we do this much better? Because we do it hundreds and thousands of times better than more than our clients. So, our ability to capture those digital traces – remember when we did that, the digital human labor, you came to us and you said, oh, you did this at 100 clients. Why don't you do it for us? We are using the same philosophy to get digital labor to be a part of our integral ecosystem.

Token spend is an architectural problem as well. Context quality, you're going to hear about context a lot today. Models don't know which context matters, the process, whatever you give them. So, you have the ability to filter the context with quality. Model routing. Not every step in an agentic flow requires frontier reasoning.

Continual learning, enterprise AI system shouldn't solve the same class of problems from scratch every time you actually get to them. Harness design, all of that put together, especially because operations work is multi-step, you can actually capture these digital traces, integrated with human effort and create a new craft which will make token economics an important part of your value chain for companies like Cognizant to not just help you externally as a partner on human labor, but also on digital labor and the integration of human and digital labor.

So, here is what we are looking at as the opportunity. It's a $6 trillion market, not $1 trillion market. And we'll come to the $1 trillion market on systems which we have been working on for the last 25, 30 years. Software engineering, integrated with agentic into business operations, business flow reinvention, business flow reimagination, enterprise AI guardrails, context engineering. We spoke about context engineering since 2024. It's the new craft. Writing code was a craft we built. Assembling context is going to be a new craft we are actually developing and we have a partner with us who we work with, you're going to see in the later part of my presentation.

Token harness, agent harness, workforce, human workforce and digital workforce integration, outcome-based services. That is going to be the mix in the future of what companies like Cognizant, we call them AI builders, will do in the future. And that's the reason why we think we can bridge that gap between production value to enterprises and AI capability, which is continuing to go up and up.

The bridge to production value graph is a three-vector strategy. I have spoken about it since almost 2024. Vector 1 is software engineering, rebuilding – I mean classical software was done in a way, we are going to now and in fact as we talk we have made all our software engineering autonomous, asynchronous agentic with humans in the middle of the chain as software engineering.

In the last two years, we used that opportunity to create more for less and the paradox of living for more for less gave us the opportunity to push the backlog of our clients at a lower cost, higher velocity. It also gave us opportunities to consolidate. In fact, every quarter we did a large number of $100 million deals based on consolidation, but it also created throughput for increased software spend. Every time – this has happened in the last 50 years, when technology has got cheaper and technology has got easier, more technology got consumed. And for a moment, even if that elasticity doesn't exist, you still know that there is a $5 trillion of new spend in operations of companies, which is getting unlocked because systems and people are now getting blurred and systems and people are – now systems and people and digital labor.

So, we see the Vector 1 opportunity by itself on its own, standing for a higher velocity and a higher elasticity of spend. And I'm going to talk about some of the value pools. Vector 2 is all about industrializing AI, which is integrating AI into enterprise landscapes. Unlike in the past where we took the old technology and junked it, this is

a technology which is going to layer on top of the old technology. You could argue whether the older technology will become a system of record or it will be less consequential, but it is not going to be taken away. It's going to sit on – the new technology is going to sit on top of the old technology and integrating it, securing it, building the data funnels and integrating it with the system of record and creating the interplay between the classical software and the agentic software is Vector 2.

And Vector 3 is this sprawl into operations of companies which we believe is 20% to 40% of spend of enterprisers versus the tech spend of enterprises in the past, which was only 0% to 10%. So, this is how the production value and the gap is going to be fixed as we go forward.

In a simplistic way I see it in two swim lanes, old things done in new ways and new things done in new ways. Old things done in new ways is classical software with higher productivity, autonomous, asynchronous where new things in new ways is this agentic sprawl and business operations and building new products and new services. I mean, we have not got to it yet, right? The $20 trillion of economic output I spoke about by 2030, one-third of it is productivity. That is what we are actually now seeing in AI. As you go forward, you're going to see new products, new services and when that happens, you're going to see a very different kind of growth imperatives for companies attached to AI and consumer-led – AI product-led consumer consumption, which is going to come into picture.

So, here are the value pools. It's a little busy slide. The first block is value pools of old things in new ways. We spoke about autonomous software engineering. In addition to that, this is something I'm seeing today. This is not a futuristic thing, today. And this list will evolve as we go forward because this is running at such high pace, if you – if I do this presentation in three months, I'm going to show you a newer list.

Now, let me quickly run it through. And today for the rest of the day, you're going to see this value pools highlighted by my colleagues in more detail. First, secure AI services. It's a Y2K moment for me. Imagine doing vulnerability discovery using Mythos or GPT-5.5, and you do that vulnerability discovery at machine speed, that's not the opportunity. The opportunity is to patch, refactor, remediate these applications. And that's why I call it the Y2K moment. When the Y2K happened, we opened the lid to fix the bates and then we found a whole bunch of things, architectural deficiencies, brittle logic, and we tried to fix it.

That is what we are going to see. It's a 10x opportunity. And people – clients will be paranoid because that vulnerability is exposed. So, we have partnered with CrowdStrike, Palo Alto and Zscaler with our security services, and we have got the entire company's service lines to rally behind it. Modernization, mainframe modernization, I have already started winning deals.

We have a huge pipeline on modernization, 800 billion lines of code of mainframes, lines of code are written, and they've been in enterprise landscapes as – some of it as legacy debt for more than 50 years. It used to cost us $10 per – every line to refactor. It now costs $1.5 to refactor. SAP HANA migrations, SAP changed the dates every year because the ROI to do that was lacking by enterprises. Here we are now, we could do it faster, quicker, cheaper. SaaS reimagination, there's going to be some SaaS reimagination. Either you're going to apply AI on top of it, or you're going to take some of that deterministic logic and push it back into SaaS.

Autonomous infrastructure using AI. In fact, it's actually one of my blockbuster offerings, autonomous – running high-touch autonomous infrastructure, even for the old stuff, leave alone building AI infrastructure for the new stuff. It's actually a blockbuster offering. In fact, as a company we are already doing double digit growth in that space. Cognitive agentic commerce, based on desire and intent. SaaS-enabled AI, now AI has this unique thing to sit on top of SaaS or go on the side and you actually do agentic work and open the SaaS layer and make it

headless so that AI can access the same set of rules. By the way, we did that for TriZetto last week. We made TriZetto headless. $0.5 trillion of claims go through TriZetto platforms. The same set of clients have $4 trillion of administrative work sitting outside of TriZetto.

So, we've made that headless so that the agentic you do, you can do it on top of TriZetto, you'd be super happy. If you want to do it outside of it, you can access the same guardrails, the same security, the same trust layers which we built in TriZetto. That's going to happen on all SaaS platforms.

I spoke about context engineering. This is a contextual science. The new code is actually to assemble context. We've spoken about it since 2024 and it's now getting to reality. You're going to hear some of our clients talk about it. Sysco Foods actually implemented this.

What is context engineering? We took all the deterministic parts and wrote classical software, codified it. We took all the probabilistic parts and we put it in a human endeavor. Now we are saying this is a technology, if you ground it well [ph] into the hustle (00:25:10) of the company, into the heterogeneity of the company, you can actually create a science around it. That's the new code we are writing. Model engineering. LLM versus small language models. Sometimes small language models are more superior. In fact, my data analytics teams are working on small language models for some of our clients.

Domain specific models. Digital twins replicate what a human does and then pluck things out of the digital twin and create an equivalent of you and then integrate it. Physical AI, it's not taken off but it's going to take off. It's a much, much bigger opportunity because the leapfrog for digital enhancement into physical, robotics, autonomous systems, mobility is going to happen. We have now invested heavily to seize that opportunity as it comes. AI infrastructure, earlier this week we actually put a press release on our AI Factory and this year we are Dell and NVIDIA's AI Velocity Partner for building AI Factory.

Trust and regulatory guardrails. Agent development cycles are going to be very different. If you're building agent development cycles for new products, they're very different to the software development cycles. Agent development cycles, you scope work – you scope for outcomes, you design for behavior, you do iterative learning loops to ground the agent and then you supervise agents, you don't keep the lights on. So, ADLC is a much bigger service.

AI data training services, Cognizant's BPO arm, which is called the Intuitive Operations, has been doing machine learning data training for the last 10 years on the big tech companies. We are taking that capacity and repurposed it for AI training – model training and that's a huge service we are excited about.

So, what is the reforge of first principles we need to do to be that company which bridges the gap between production value and AI capability. The reforge of the first principles are we go from a system integrator to a AI builder, and I'll come to it in a minute, what a AI builder means. It's a bespoke opportunity because this is not the deterministic piece, this is the probabilistic piece. So, it's a bespoke opportunity, much, much bigger. And the capability is very different.

Pyramid to interdisciplinary talent. We thrived on a pyramid of talent. You're going to hear from Jatin on how that is going to spread, you're going to hear from Kathy, our CHRO, on how that mix is changing. Interdisciplinary talent at the intersection of domain and technology, at the intersection of operations and technology. That's what we are trying to put, and then embed digital labor into it.

Platform-led agentic delivery. Earlier this year we launched the platforms unit. Our thesis is our clients are going to consume this, not as services, not as software, but as platforms. The pioneering move forward looking clients are going to consume this as, give me the outcomes through a software plus people endeavor. So, we have opportunity to be a platforms plus peoples company and we used to be – we've been the platforms company in healthcare for many years. And we need to have the courage to own business outcomes because no longer are we building systems, we are building – we've blurred the line between systems and people, systems and operations. So, this makes a business more durable, sticky, more value capture, a different risk profile and a different return profile.

What does it take to be an AI builder? What are the blocks – building blocks? Full stack bespoke capabilities, integrating human and digital labor into software engineering and business operations. Flatter organization structure. This is the transformation we have to go through. We think the people in the middle are player coaches. The nodes are going to be more real time. So, the people who actually were moving information up and down, coordination, measurers, orchestrators, we shrink that, we create more flatter structure. In fact, I've been a big proponent that the pyramid is going to look much more broader. We have hired more school graduates this year – we are hiring more school graduates this year than last year. Last year we had 20,000. The year before we had 12,000. So, this is a more profitable business. If we can push the pyramid broader.

New roles and new job families, which you're going to hear from Kathy, and lots of jobs on the front and lots of jobs on the back while AI is in the middle. There is a ton of job families and jobs which we have now established in this flow cycle.

Business process reimagination and reinvention is the foundational step. We have a BASIS methodology which is driven by our consulting teams trademarked by us for transforming existing flows. We spoke about context engineering, the ability to put the work graphs, the guardrails, the sensors, the tribal knowledge and capture that, assembling that context is the new code.

Somebody asked me this question, what does this new Frontier Engineer do? And I'm going to come to that in a bit. And I'm going to talk about it. And that's an important part of this process.

Our partnerships are now layered with frontier models, more AI-native companies, in addition to the incumbent SaaS and security companies, and the hyperscalers we worked with.

One of the craft Cognizant built over the years is to take any new capability and build it at scale. We don't buy capacity. We build at scale, if at all we buy we buy it for beachheading it. What is a Frontier Engineering and a Frontier Operator? I mean, the market a little bit gets confused and they kind of use the Forward Deployed Engineer as a proxy to it. The Forward Deployed Engineer was built at a time when the world problem was everything in a company is broken, I'm going to stitch everything together, bring it together on an ontology, I'm going to bring a black box and the engineer who will fix it. That isn't the world problem we're trying to solve today.

The world problem we're trying to solve today is you have flows in a company. You have to open the flows, audit it, do the evals, integrate it back and reinvent the flows and deploy it? A Frontier Engineer is somebody who brings that machine, it could be any frontier model. It could be different frontier models for different tasks, that's the craft we have and you reinvent those flows. So, you need to have capability at the intersection of technology and a domain. A Frontier Operator, some of our clients will say, you know what, if this has digital labor and human labor, I'm not going to run with this. My core mission, if I'm a healthcare company, is to underwrite – underwrite my insurance and to acquire clients. You manage the operations.

We need a BPO plus-plus-plus which can use digital labor and human labor together to deliver those outcomes. That is the kind of capability we're going to build and we're going to build it at scale. A Frontier Engineer and a Frontier Operator, we're going to build it at scale so that all our clients can actually thrive on this. Flatter, leaner, modular squads, non-STEM disciplines because we are now getting into other areas.

We have new skills – new ways of skilling and new ways of learning. You're going to hear from our Chief Learning Officer about our Skillspring training platform, which is actually a micro personalized, at your pace, and it navigates to the finishing line. We have an AI fluency meter for the 350,000 employees. We're flipping this and now – now we are actually presenting it to our clients, embedding it into the work we do. In fact, we believe learning is a part of our infrastructure stack – our AI infrastructure stack. So, that's the change we are making inside the company as well as the transformation we do with our clients.

Our platforms business is a new business unit. We have a heritage with TriZetto, we have 200 million lives on TriZetto. We want to make that a 60-rule business. A 60-rule is something which margins plus growth kind of gets there. That's our aspiration. We sense what we need from our labs – our award-winning AI labs, you're going to hear from my colleague, we build/buy/partner/invest. We have a new investment arm and we think we can do it in two swim lanes, engineering to increase velocity and cost of deployment of AI and business operations. So, business platforms like TriZetto, either we build, either we buy, we invest or we partner. Because there are – I mean, look, software companies have to reforge as well because they have to own outcomes. So, they will be looking for partners who can actually take the end-to-end responsibility. We have physical AI, which my colleague is going to talk about, and we are building a platform called Intelligent Spine.

So, here is the shifting economics from labor to outcomes, token optimization. AI-infused rate cards. A few of my clients are already using the AI-infused rate cards, A0 being all human effort, A1 being human effort audited by machines, A2 being machine effort audited by humans, A3 being autonomous. So, we're embedding human and digital labor in the rate card, because some of our clients are saying, I can't fix price this, I can't do managed services, I can't give it on outcomes because we have products between your teams and my teams. So, we have AI-infused rate cards, I have clients who are already using it. And we are already underwriting outcomes in TriZetto because we do it per person – per member, per month for TriZetto.

So, in summary, the bridge from AI capability to production value needs a company like Cognizant more than ever before. The opportunity is $6 trillion of market versus $1 trillion of market for systems to a – to $1 trillion – $6 trillion of systems plus people, plus digital labor. And the capability gap, the production value to capability gap, you need to be AI builder and Cognizant is poised to be that AI builder that helps enterprises bridge that gap between AI capability and production value.

So, thank you so much for listening to me today. What I'm going to do in my next section of this – today is to invite a partner of ours. The company is called Workfabric, Rohan, the Co-Founder is going to be here. We've been working together on context engineering since 2024. We have done some pioneering work for our clients. These are a bunch of researchers from Harvard University and Carnegie Mellon who put this up. And we work together. We've also applied this to our own selves. And he's going to tell you a little bit of context engineering and also show you a little bit about our sales transformation and how we're using context engineering to generate a pipeline. And he is also going to show you about our customer success program, which is about anything which happens in a client we sense it, and we tell our clients before our clients tell us.

So, these days, if any of my clients wants to call me up for a potential friction, I should be able to tell them before they tell me, and I should be able to tell them how I'm addressing you. So, we are able to do that with the power of

context engineering. And he is going to show you some of it. And we have powered our sales opportunity management system using context engineering.

So, Rohan, over to you.

Rohan Narayana Murty

Co-Founder & Chief Executive Officer, Workfabric, Inc.

All right. While we just wait for this to load up, the context to what I will show you is, as Ravi mentioned, in October 2024, Ravi and I met and at the time the predominant world view was that, as you saw in Ravi's graph model capabilities will keep increasing, model intelligence will increase. But the problem with that is models don't really understand organizations, the nuances, the specificities of organizations. So, how do you actually solve for that? Because unless you solve for that, models will never speak the language of the organization and produce value in the context of the organization.

At the time, this was not a popular idea, certainly not an accepted idea. It was just two guys in a room who believed this idea, zoom to a few months later, we had an actual working deployment. We had built technology that could learn from how people worked in organizations that could extract tribal knowledge inside organizations. We had deployed this in a Fortune 500 company. We saw some results and the transformation in terms of what agents could do when you grounded them in the reality of a company in the context of a company was dramatic. The outcomes were dramatic.

Ravi and I published our first article in Harvard Business Review in April of 2025. Insofar as we know, it is the first ever documented case of the value of context inside the enterprise. Since then, this whole space has taken off. You'll pretty much hear everybody talk of context engineering. What I'm going to show you today, though, is a live demonstration of what Cognizant and our company are doing together with context engineering. I'm going to show you outcomes that, at least as far as we know, have never been possible before, nor have we demonstrated anywhere else.

Okay. All right. So with that, what you see here is a live system. This is a twin of various accounts inside Cognizant. Obviously, for all the right reasons, we've had to strip out names and logos, but the data is very real. Okay. And the use case here is as follows. Any company, certainly Cognizant included, or any company that is selling to another enterprise, has a reasonably large surface area that interfaces or touches their customers. Sales talks to customers, delivery talks to customers, support talks to customers, finance talks to customers, et cetera., et cetera, et cetera and so on. But the intelligence about what is happening when you interface with that customer, what you're learning about that customer is not captured in any system of record.

In fact, your CRM pretty much only captures what your salespeople put into it, right? But what about all of this other intelligence? Maybe there's a junior person in your organization who learned something very relevant about your customer that could have been useful to your salespeople. Maybe your delivery person learned something and so on and so forth.

So, there is this incredibly powerful intelligence about customers that organizations are always learning. But this intelligence is living in silos beyond the sales teams, but historically used to thinking of having a record – system of record like a CRM, and the system will tell you the truth. Well, that truth is only good as whatever you put into it. But what about what people are learning, the tribal knowledge, the tacit knowledge that you're learning about your customers? And so what we have done together with Cognizant, in our opinion, this is the first of its kind use case in sales is we have used context from within and across Cognizant to find net new revenue opportunities that no human could have ever found on their own.

Okay. And what you see here is each salesperson in Cognizant when they log in every day, they see for the accounts that they're mapped to, leads that AI is suggesting that they should be chasing. And these leads are generated using context from across the entire organization versus only what is in sales. And therefore, by definition, humans can't find them. And for a moment, if you think about what I'm saying, it's really important because the most dominant use case of value for organizations from AI has been code generation, or at least that's what's widely accepted. But what we are talking about here together is this is direct ROI because we're producing net new pipeline.

I'll spare you the specificities, but I'll give you a sense of what is possible. So here, as an example, it's telling the salesperson you should go to Apex Auto – obviously not a real name in this particular case, the story is real – and go pitch Cognizant's Flowsource for QA optimization. And now when the salesperson clicks on this, what this is showing him or her is based on the context learnt across different teams in Cognizant, there is the staffing levels for QA at this particular customer are under scrutiny.

Now when we double click, there are more details because we are capturing this context from the lived reality of these teams, from emails, from contracts, from documents, from conversations et cetera, et cetera and so on off Cognizant's own people. While making sure we have privacy and customer sensitivity and all of these guardrails put in place, we're learning from the sales teams, from the delivery teams, from people across the globe that there are signals that are telling us that in this particular customer there may be some increasing oversight on QA staffing levels.

And the second point is, looks like this customer has a mandate to reduce engineering costs by 15%. And the way we actually learnt that was by listening or by learning from a couple of people who are spread across the globe who had picked up various bits of information, that this was a stated objective for a customer. Some of it was in some ticket that was seen in Manila, some of it was seen in India, some of it was in the US and so on and so forth. And when we put these signals together, we proactively realized what the priority of this customer is, what is their pain point, and therefore what can Cognizant proactively go do for this customer versus wait for them to come and ask you something.

And so therefore in this particular case, it's telling the salesperson, go to this particular customer and pitch. Here's a proactive pitch. The customer didn't ask for this, but go proactively pitch for them saying, I think you may have this pain point. Here is how I can be of service and of value to you. And then of course, there's a briefing generated, an email to make it for the salesperson and so on.

Now what this is essentially doing is it's producing a class of revenue opportunities that are synthesized by the living reality of what Cognizant's own employees are learning about their customers, right? Now, later on I'll – Ravi will tell you there's a dramatic number because this is not a pilot. This is not a proof of concept. This is not an experiment. This is not a research project. This is a living, breathing system that is running on thousands of Cognizant's people, their teams, their entire global sales force producing this kind of an outcome. I'll let Ravi give you the dramatic big number.

Now the second quick thing is – all right, so now that we have context, you take an account, you have context of the account from people who touch the account from these different teams. And so Ravi mentioned, he wants to be able to go to a customer and say, here is a problem that I'm already aware of and we are fixing it versus wait for the customer to come in, escalate to you and say, there's a problem, I'm not happy. So the way we address that using context is now we have built these mission-driven twins of these accounts.

What does that mean? Well, so we take the context of all people who work on an account. We put that together and we've created a twin of the account itself using this context. And now we give each one of these account twins, a mission. So, you're seeing – again, these are real accounts, names have changed, obviously, but each one of them have different missions. So, for example, there is an account where the mission is prevent revenue churn. There's another account, where the mission is find more revenue that is likely to close in 120 days. And these missions are live. So, when Cognizant is not looking, it's supposed to be doing its job and finding new revenue that people may not be aware of or build a renewal strategy and so on, so forth. You can pretty much create any kind of mission.

But I'm going to show you one particular account and the mission there, again, the facts have changed very deliberately, but it'll give you the flavor. So, this one needs attention. Okay. And this is Cognizant proactively recognizing that it needs to pay attention to this account. So, let's click on this.

Again, remember, this account is collecting context globally from everyone servicing this account. So, it turns out, for example, on May 11, this mission was accepted by this twin. All of this using context, in this account they were doing – Ravi earlier alluded to modernization and SAP migration, this account is doing SAP migration in multiple sites, now based on the context of the various teams, looks like the Dallas site went fine, the Sydney site seems to go fine as well, slight anomaly in Riyadh, but that's okay. It's a holiday.

But then there's a first sign or a signal that is telling us saying, there may be an issue in Sao Paulo. And we picked that up because somebody within Cognizant learned something from the customer, giving them some kind of feedback. Usually that would have just been lost or you would have waited until that accumulated or escalated.

And eventually the signals get stronger. And it looks like the deployment in Sao Paulo. And to give you a sense, the context is from these various teams, so the onsite team, the delivery teams, et cetera, et cetera and so on. And there are various signals that are embedded in these teams that are telling us, saying this particular site may have an issue.

Again, now post this, this account twin pulls context from the SAP Practice twin, there's a twin even of the SAP Practice in Cognizant saying is what's happening in Sao Paulo based on its context very different from what you usually see. And if it is, okay, it looks like we need to do something here. It suggests to the human, maybe you should pause this now – the human takes the decision, so ultimately the humans are still in control. Pause this and then the twin comes back and says, here are two people, based on their own context who might be experts in helping resolve this situation.

Eventually it also helps draft an email to the customer and so on, and then it creates a briefing for the CEO because in case there's an escalation, or if the CEO wants to pick up the phone and call them saying, don't worry, we know there's an issue with the Sao Paulo deployment, but we will take care of it. It's a very detailed briefing. All of this, the key point is powered by the tribal knowledge and the context of individuals inside Cognizant.

Okay. And there's obviously more to see and more to show you. But hopefully I've given you a sense and a flavor. You hear the term context engineering, but I wanted you to see – we all wanted you to see what Cognizant is actually doing in production with real numbers and real deployments. And hopefully this gives you a sense.

And with that, I'll hand it back to Ravi.

Ravi Kumar Singisetti
Chief Executive Officer & Director, Cognizant Technology Solutions Corp.

So at this point of time, we roughly have $200 million of pipeline generated incrementally through this extraordinary effort of doing a sprawl on the systems, emails, meeting, chats, everything else and generating it. By the end of the year we think this is going to be $1 billion pipeline just by listening to conversations of our clients and extracting the tribal knowledge to create some insights.

In fact, there's another interesting use case which Rohan didn't mention, which is about how do we staff programs. Today we staff based on resumes and skills in a system, every time we have to staff something, we just have to ask the context engine. Tell us who in the company is doing similar work? Who will fit the bill? It'll find people around and say, these are the 10 people who can fit the bill. So, we've done that as well

With this what I'm going to do – now do is everything we spoke about. What does it do to our operating model? What does it do to our business? Jatin and Ravi Kiran are going to talk about it. And Jatin is our Chief Financial Officer. Ravi Kiran is our Chief Strategy Officer. And before they come on stage, we're going to have one of our customers, LPL Financial, which is actually working with us to provide the wealth management platform with agentic to support their wealth advisors. So, we're going to see the video and then get to Jatin.

Nitesh Ambastha

Executive Vice President and Group CIO-Firmwide AI, Digital Platforms & Investor Experience, LPL Financial

In wealth management, the scarcest resource isn't capital or talent, it's time, specifically the time an advisor gets with a client. Nearly a third of advisors say, they don't have enough of it. That's not just a productivity issue, it's structural, and exactly the kind of problem AI is positioned to solve.

I am Nitesh Ambastha, Firmwide AI leader at LPL Financial. We serve more than 32,000 advisors across independent and institutional [ph] marks (00:53:22). Many operate as entrepreneurs who own their practice and client relationships, while others operate within institutions. Across both segments they look to us, LPL, for the infrastructure to compete and deliver for their clients. So, when we think about AI, we are not thinking about it as a feature to ship, we are thinking about AI as a multiplier of human judgment at the scale of tens of thousands of advisors and institutions.

I leave you with three ideas. The first, AI doesn't replace the advisor, it multiplies them. The future of our industry isn't human versus machines. It's a compound model, a financial professional and AI working together where the human owns the judgment and the relationship and the AI removes friction from everything else.

The second, generic AI produces generic outcomes. Wealth management is highly regulated and deeply personal, with advisors operating under fiduciary and best interest standards. We need AI that is secure, compliant and designed with clear data guardrails. In our industry AI has to be built for purpose.

The third and most relevant to this room, the next five years of wealth management won't be defined by raw models scale alone. It will be defined by who has the right partnerships. At LPL Financial, we take a deliberately modular approach, curating and orchestrating across a focused set of partners where there is real depth in R&D and domain expertise.

Cognizant's investment in its Neuro AI platform and the financial services specific capabilities they are building on top of it reflects the kind of focus we look for. We are working together on a targeted set of initiatives and you will hear more in the coming months.

So to the analysts and investors watching, when we evaluate partners, we are not looking for the loudest AI story, we are looking for sustained R&D, industry specific playbooks and execution. When those come together the

impact for an enterprise like ours, for our advisors and for the millions of clients they serve can be outsized. Thank you.

Jatin Pravinchandra Dalal

Chief Financial Officer, Cognizant Technology Solutions Corp.

Hi. Good morning, everyone, and thank you for being here. Very happy to see so many of you being present to hear our story. The purpose of my and Ravi's session is really to speak about the journey for our customers from something that look incredibly exciting from technology potential standpoint to real value creation for themselves, and therefore their own shareholders, right? I mean, that is the end goal. And for that, you need two things. You need a great strategy. And I hope you heard that well from our CEO. And I'm going to speak about how does the company behind changes to align with this strategy and deliver that value to our end customer.

We think this operating model has five legs, which are crucial. The first is how we sell. The second is what we sell. Third is how we deliver the new offering. And fourth, which I'm sure many of you are very interested in knowing, how do we monetize this expanded surface area? And finally, how does the talent or the organization that all of you have seen touched, understood for last three decade changes and becomes an organization of future. These are the five elements of our operating model which are very crucial for the execution which one of our customers just spoke about, that execution has to come from the operating model, once you know what the blueprint is.

We'll start with GTM first and I'll request Ravi to speak about it.

Ravi Kiran Kuchibhotla

Chief Strategy Officer, Cognizant Technology Solutions Corp.

Hi. I think from a go-to-market and the way we take to clients in terms of how we sell our services, obviously every technology shift that has ever happened always has led to IT services players like us helping clients through their journey, that has always been the case. And when you look back at the cloud era as an example, I think we've been able to help our clients with the entire lift and shift to the cloud, build native-cloud applications, do the integration. A lot of the infrastructure that was needed underneath, cloud, was kind of something that we have navigated our – most of our clients through that journey.

I think with AI it's going to be no different. I think AI – in fact what Ravi alluded to, being an AI builder and the fact that there'll be more and more bespoke AI capabilities that have to be built needs a massive uplift from service players like us. And I think it's also the complexity of bringing and orchestrating a whole bunch of ecosystem players. So, you're looking at frontier model players, you're looking at cloud, you're looking at security and now you're looking at infrastructure, a lot of things that have to come together.

So in that sense, I think it is going to be a massive need for players like us to step up the game and take it to clients. One of the things, obviously, from an operational standpoint, that we have enabled is creating AI market unit, making sure that we're bringing the strength of the organization and bringing it as a single interface to our clients. So, it's something that's critical to taking it forward.

The second is deepening our partnerships with many of our partners, the frontier model players. So, one of them – some of the security players that Ravi alluded to as well, there's a lot of that in terms of how we kind of bring together a far more integrated model that we can take to clients where there's more of token and inference models that can embed into our go-to-market strategy itself.

And the last thing is obviously about – you've seen the demo from Rohan. I think it is a phenomenal demonstration of how we have leveraged context engineering to kind of build our sales capabilities. Our entire sales engine is enabled with these context sensors and the capability behind it. And to me I think this has only amplified the sales capability to take it to the next level. And there's more – obviously we're seeing a lot of benefits that are coming out of it. But to me, this is, in a quick summary, the operationalization of some of our GTM activities.

Jatin Pravinchandra Dalal

Chief Financial Officer, Cognizant Technology Solutions Corp.

I want to speak about what we sell and Ravi had a whole slide around Vector 2, Vector 3, and I would just double click on one or two of them, how this is changing. And, again, I mean, we all have been in the industry for more than two decades, almost everybody in this room. You know this, that whenever we provision something that goes first and security follows, and then you put a wrapper of security around it. And if you are really smart, you think about security even as you're developing that.

Now, look at how AI is changing this. And I'm giving this example because when we speak, the first thought that comes to our mind is – but AI is deflationary for services. But I'm telling you an example of a service that is in force because the AI exists, which is Mythos, you start first with security. And look at this, the entire paradigm of three decades has changed. Security was following, now security leads. Security at the front tells you that these are the vulnerability which is sitting in your organization. Now go fix it.

So, this didn't exist six months back. This exists today. I'll tell you one thing else. If you look at in security, every time you put an element in your network or element in your outcome that you are generating, you need an extra element, the database, network, applications, cloud. You know that your earlier security posture has completely changed. Now you need to re-secure it with a new element of AI which is now helping you deliver services. Very easy to say okay, let's put AI and get a better outcome. But you are not thinking what you know differently on security to – just on a business of security, you've now two very large opportunity.

The first, what we call is AI for security. And second is secured AI, by itself large momentous opportunity – Y2K opportunity for us. Now, if you look at the second is really operations and tech. Traditionally, operations and tech have run parallelly. IT services companies have large operations practice. But even when customers make a decision, they said, look, let's give operations to player A, tech to player B. Now you have agents which are sitting in the middle where somebody can do what client is doing today, which is teaching tech and ops together and create an outcome. That's the whole new opportunity.

One simple example of a new service, which has come because of AI, is this training the agents, which we do in ops, which was not present again a year back. So, we feel very excited. The composition of services that we are going to sell is going to change materially in times to come. And even our inorganic investments are all aligned on that. 3Cloud that we acquired in the beginning of this year, the leadership here, I would encourage you to speak with them, but they are really focused on supporting organization in their cloud journey as they leverage AI.

Astreya really focused on creating – on helping create the data center – the technical help to create data center which we all know is a big investment in today's time. So, we feel very energized about the new offering and the new pools that we can tap with AI.

Ravi Kiran Kuchibhotla

Chief Strategy Officer, Cognizant Technology Solutions Corp.

Okay. The third dimension of our strategy and the operating model is about the platform shift. Unlike the past, where it was all about skills and platforms, we are now shifting into an era where it's going to be skills, agents. You're going to have a fairly large agentic layer and AI-infused platforms that are going to co-exist going forward. So, in that sense, I think a clear shift towards platforms that are going to be more outcome-oriented, they're going to actually start measuring real value. And one of the operating principles that we have used behind it is to start looking the set of platforms that we have, and we've consolidated all the platforms into one unit. One of the first things that we have done.

The second is obviously how do we look at AI enabling [indiscernible] (01:04:48) even going to the extent of making them AI-native and coming up with new monetization models so that the platforms can actually deliver value as part of the offerings itself to our clients. And, in a sense, I think the entire platform approach something that we think is also important from a nonlinearity standpoint, eventually we want to get to a stage where the workforce is complemented with the agents and the platforms so that there's a fair bit of nonlinear economics that we can derive out of it.

Jatin Pravinchandra Dalal

Chief Financial Officer, Cognizant Technology Solutions Corp.

The monetization is crucial. Ravi we spoke about headless, TriZetto. Now imagine, I mean, we always have sold TriZetto in a per member, per month basis and that's a real constraint as only so many new users you can get in organization. Now that IP is exposed to the whole agentic workforce of the world. Anybody can come in tap that IP and we will have to monetize – we have an opportunity to monetize that interface.

We spoke about A0 to A3 inference building inside the way we sell, and you could argue, but inference could be margin dilutive. And I would say no. The reason for that is that, as Ravi mentioned, there is a secret sauce about how we use inference. Just the way we have a pyramid of human skills on left side we'll have to build a pyramid of inference on right side. And then we leverage that pyramid to build the best outcome for our customers. And that's the secret sauce that we can go with. So, now traditionally we have sold skills, now we sell skills plus inference. That's massively large opportunity that we never had before. And that's the $6 trillion that Ravi spoke about. And this is not something tomorrow. This is, even as we speak, not a $5 million deal, I mean, $6 million deal, these are large deals that we are submitting, which is built up of skills plus inference, and this is the outcome that we are generating.

Ravi Kiran Kuchibhotla

Chief Strategy Officer, Cognizant Technology Solutions Corp.

We talked about how we sell, what we sell, how we deliver and how we commercialize. Now, bringing all these things together, obviously the foundation to all this is the workforce. How do we retool and reskill the workforce? That's going to be one of the foundational shifts that we are doing. I think Ravi talked about the new roles and the new career architecture that we have launched, the Frontier Certified Engineer and the Frontier Business Operator. These are new roles that are evolving for this new AI era, where they're going to be a deterministic and probabilistic ecosystem that's going to co-exist. That's going to be a big part of the workforce reskilling that we are currently underway.

Obviously, when you look at the way AI is going to impact a lot of our clients, the pyramid structure that we currently have today is almost no longer relevant. I think we're going to be a much flatter organization. It's going to be a much wider organization. We're going to have a lot more agents at the bottom, a lot more interdisciplinary skill sets that are needed for the future. So in that sense, the middle layer essentially will be doing more of a player coach kind of thing, that kind of Ravi alluded to earlier.

So, this is a fundamental shift that we are doing and the AI fluency is something that is a critical part of our strategy as well in terms of bringing together not just how you get trained, the kind of project exposure that you have, the kind of whiteboarding and other related hackathons that you're part of in various forums. How do you bring all these together to provide a fluency score for every individual? And eventually, how do we take the fluency score concept to our clients as well? So, how do you make the client also kind of? So, this is where we are headed from a workforce transformation standpoint.

Jatin Pravinchandra Dalal

Chief Financial Officer, Cognizant Technology Solutions Corp.

If you want to really takeaway three things that keep us driven every day on our operating model. This is this, convergence of tech and ops and the new opportunity that comes with it. Blended effort. I mean, we no longer think man-months, we think about the total blended efforts of skills plus inference plus platform. And finally, this two combined gives us the larger commercial area and how are we monetizing it. Because this is really the key of keeping 350,000 employees grounded on where we want to go and how we get there. And there are – of course, there are KPIs. Some of them are old world KPIs which will never go back to like man-months growth.

Some of them are transient KPIs like token consumption. I mean, we were very excited about it six months back. Now we think it's no longer relevant. What I should measure is the throughput that I generate out of tokens. And some of them would be now in external domain like revenue per employee, operating profit per employee, machine-assisted code and some of them will remain internal as we transform the company. They will be relevant for a period of time and then they won't be relevant anymore.

This is a quick summary from Ravi and me on our operating model. Thank you very much again for coming here. People who are listening our webcast. Thank you very much for joining us. Our IR team and our media team is on standby if you've any follow-up questions. But, thank you very much.

Ravi Kiran Kuchibhotla

Chief Strategy Officer, Cognizant Technology Solutions Corp.

Thank you.