

R-Srikrishna-CEO-AI-day-Session-1

We'll get started. Good morning, and welcome to all of you. Now, thank you for spending your day with us.

I want to begin by saying that we feel very good about where we are as a business. And we feel good because we're winning deals, whether it's in existing clients or new clients. And now, for some time, the majority of our deals are won or lost because of AI capabilities. They're infused with AI, or in some cases, they are only AI. But virtually everything is infused with AI. And we're winning these deals because we think we have differentiated capabilities. We think we are best in class. And this is agnostic to size. I'm talking about you could compare us with an Accenture or anyone else. We think we have best in class capabilities. So the next several hours, we are hoping to give you a glimpse of these capabilities and why we're winning. We're going to try and bring it to life, both through describing what we do, but also some customer examples.

You know, when you all walked in, we would have given you a bag with this book. It's a customer impact series on AI. We published the first version in April. That book was about one third the size. In three odd months, the number of case studies, the new case studies that were published, has gone up more than 2x. So, this is an indication of the velocity at which both our customers are adopting and they're able to execute for them.

Now, before we talk capabilities, what are the strategy underpinning the capability? Our strategy is designed for two outcomes. The first outcome is to create new avenues of growth with AI. Now, deflation is real. We're not running away from it. In fact, we are being proactive in addressing pockets of deflation with our clients. But we think there is a massive set of opportunities to outrun this deflation and sustain high growth. In fact, in our case, get back to high growth. The second objective is for us to improve profitability. Clearly, AI has an opportunity to create non-linearity and with it is an opportunity to improve profitability. Our first objective, though, is the first one, which is growth. I think profitability is not our day one objective. In some cases, it could happen, but the focus on day one is momentum, market share gain, and growth. So how does this translate for customers? For customers, at the highest level, we take services to them in three blocks. The first block is AI for IT. What we're doing in AI for IT is to infuse AI into everything we do, but also several new lanes or services that become possible as a consequence of AI. The second lane is AI for business. And AI for business means something slightly different for us than it does for many of our peers.

So I will talk to what it means for us in a bit. But at the highest level, AI for business certainly includes impacting the core business in a material way for our customers. Now, both these have to be underpinned by a layer of trust because AI adoption at scale cannot happen without a foundational layer of trust. So these are the three broad things that we take to market. Now, if our strategy is to outrun deflation materially with new sources of revenue, our first imperative is agility. We said earlier this year, beginning of this year, that we will launch one new service every month on AI. Actually, we've been doing somewhat better than that. We also said every time we launch, we will knock on 100 customers' doors in 90 days from the day we launch it. Now, for us to take one service to market, we need to create more than one service per month because not everything we create will meet all the bars to take

to market. So there is an engine for us that allows us to be agile. There's an engine that creates lots of cool new stuff, and that engine is our innovation lab. Our innovation lab mainly is in Chennai, but it's also in other locations. It has more than 700 people. When I visit them, I get inspired. I get inspired by how smart they are, how AI-native they are, the speed of how they're able to think and execute. The average age in our lab is only 27 years old, and we've equipped them with the best of everything, whether it's software or hardware, to create new stuff. Virtually everything we've brought to market in the last year is from platforms they've created, which allows us to build and take a service to customers based on a differentiated platform.

Some quick stats. There are 70 little products they've created which are in production with the clients. For many clients, for many services, we are client zero. We eat our own dog food, so 90 of them are in production inside Hexaware. And there's a much bigger pipeline, and we do have a quick session on this later today. The latest innovation they've created is a platform based on which we launched a new service earlier this week.

"Every enterprise has a North Star. The ambition is clear, but friction stands in the way. Whatever its source, Hexaware's Zero Friction Enterprise works to drive it towards zero. AI is scanning code faster than ever and uncovering more vulnerabilities than any team can keep up with. Zero vulnerability finds threats and contains threats sooner before they escalate. Legacy code doesn't just slow you down. It's where your biggest vulnerabilities live. Zero tech debt modernizes legacy architecture to keep old code from dictating the pace of change. The friction continues into delivery. A stalled backlog. Late surfacing defects. Speed of thought. Zero backlog. That's the new pace of shipping. Zero defects. That's the new standard of code quality. In production, the best ticket is the one that never lands. Zero tickets resolves more incidents before they ever reach the queue. And software dependence carries its own cost. Eliminate SaaS. Eliminate heavy license costs. Zero license replaces software dependence with AI native capability enterprises can own. None of it stands without a foundation. Infinite trust, spanning data readiness, security, governance, and observability for AI to run responsibly at enterprise scale. And the intelligence that drives it all, ZeroVity, Hexaware's AI delivery layer. One command centre for the enterprise. AI brings speed and scale. People bring judgment and accountability. When friction no longer sets the pace, the North Star does. At Hexaware, zero friction is where infinite momentum begins."

I think the power of AI is not just to, you know, improve white-collar efficiency. It is to create new things. And, you know, as an industry, including us, we've been talking to customers about what AI can do in silos. What we are shifting now is we're telling customers, think of a possibility of creating infinite momentum in your business.

And how do you do that? By removing friction in every aspect of your tech and your business. We've identified six zeros for removing friction. Those are the six pillars that we are starting with. But I can tell you this, that our imagination and creativity is the only reason why we only have six. I can assure you we will have more zero pillars with time as we progress. But first, there are six, and I want to spend a couple of minutes on each of the six pillars.

Mythos became, it's still in its kunk works, like 150 companies around the world have it. What we know from what they have is that it is discovering vulnerabilities at machine speed. Now,

even if Mythos doesn't come out of kunk, there are slightly lesser models that can maybe do 80% of what Mythos does. You're discovering vulnerabilities at machine speed. Bad actors can take advantage of it and do bad things to you at machine speed. There's no customer in the world that is ready to deal with this machine speed. What is also known is that the source of vulnerabilities is often open-source systems and legacy tech. So, solving for vulnerabilities is going to become one of the core corporate sustenance issues. And solving for the source of those vulnerabilities, tech debt is also going to go from a technology imperative to a security imperative. So that's the first two zeros. Zero vulnerability, we just launched that service, and zero tech debt. The vision is how can we get to zero tech debt. Now, I should maybe step back for a second. We're not promising zero to our clients. We are promising, except maybe in one that I will touch upon, we are promising to push the envelope to zero. We are promising a journey that is faster than what anyone else can take them towards the concept of zero. Zero tech debt, I think, is that. How can we get to as close to eliminating zero tech debt? Now, you will see in the afternoon when we take you to a lab, we have an agentic factory for cleaning tech debt. We'll show you that at work. Zero backlog. The number of clients who say, my business or the business side of the clients who say IT cannot build what I want fast enough. This is definitely true in enterprises, but it's way more true in product companies. Because for product companies, the features building is not just about improving functionality for employees. It is revenues for them.

It is market share for them. It is new market entry. And being able to build features rapidly, this is what zero backlog is. And this is actually a zero that we think we can get customers to not in a long period of time, but in six months, in nine months, we can get customers to zero backlog. There's no point building code with defects. This is one of the initial problems of AI, that a lot of bad code, too much volume. How do you ensure zero defects in the code you build? That zero defects. The best ticket, like saw the video, is the one that never came. Now, this has been a mission for IT industry, certainly a mission for me. I've been involved with this business for the longest time. How do you get an enterprise to fully autonomous, zero-ticket operations? I think the power of AI makes it more real than ever before make that a possibility. And finally, if you look at clients' IT budgets, number one is human labour, insourced or outsourced. Number two, license costs. SAS, COTS, zero license addresses that line item of budget directly. Not by implementing, that implementation is still in the human cost. I'm talking about the actual cost of licenses that customers pay. Zero license addresses that. Now, we think these six are not just independent, they're all interrelated. You build great code, you're not going to have vulnerabilities. You build, you eliminate tech debt, you're going to reduce vulnerabilities. You build great code, you're not going to have tickets in the system. You build zero tickets, you're going to reduce tickets in the system. It's a highly interconnected framework. I will talk more about this in a minute when we talk about our platforms.

But these have to all sit on a solid foundation of what we are now calling infinite trust. So infinite trust is the foundation, six pillars of zero, get you to a zero friction organization to help you realize vision of infinite momentum. Now, the three most important pillars here are first is data. If you don't trust your data, if you don't know where it came from, you can't trace the provenance, you can't govern it, AI is not going to work at scale. You have to be able to observe, monitor, govern AI agents. And you have to be able to secure them, not just the rest of the enterprise. So that's the pillars behind infinite trust. Now, I also color-coded this slide. I'm going to speak for some time about what the colour codes means and the implications for us as a company. The first colour code is red. This is deflation. So, there's a lot of chatter in

the industry about deflationary services, so this is where we think deflation is. IT outsourcing, deflation. There's no new volumes as a result of IT. There are some other pockets which we'll talk about later, is deflationary. Testing, zero defects, is deflationary. The orange is zero backlog. So that colour code is deflationary, but significant higher volume that may offset all or some of it. So right now, customers are focused on using AI to reduce the spend for software development. But the real value is going to come when people are using AI to build products and applications. That phase is going to require more applications to be built, and that's the volume increase. Everything else for us is green. Zero license I already spoke about. It's the second largest budget item in a customer's budget. Zero tech debt. For some people in the industry, there is already a big business in modernization and legacy, but for us, it's marginal. What I think, the problem has always been there of legacy code, but it's too expensive, too risky, too time-consuming to fix it. AI can solve for all of this. Our platforms can solve for all of this. So, we think this is a completely new opportunity for us. And zero vulnerability definitely is a new opportunity. Now, you look at the bottom. There's significant, all of these areas are new growth and opportunity areas.

Now, we put some numbers here for time. These are estimates. Zero license, the SaaS industry is \$900 billion. So we need to find \$3 billion to convert. We'll be good for a period of time. So, these times are estimates of what we think could be addressable and what we think could be actionable. I mean, there are estimates out there, just COBOL code conversion is a trillion dollar opportunity. We don't think that's true. We put what we think is a good estimate of an actionable time. This service, our lab created a platform called Zerovity. Now, all of us in the industry, including Hexaware, we've had different platforms for solving different problems. I have a platform that does great coding or software engineering. I have a platform for data engineering. I have a platform for IT outsourcing. I have a platform for cloud transformation. But think of an enterprise. They're in one business. They have one set of policies. They have one set of security controls. They have one set of architecture guidelines. They have one set of branding policies. Everything in an enterprise is common. How you code, how you build a server has to be based on that same common context. And that's missing. So, this is the most important. There are many other differentiated elements, but the core concept is there is a common enterprise context that then drives all of these services. Like I said, these are not individual services. They're all interlinked. This interlinkage comes to life because it's a common enterprise context on which we can deliver multiple services.

So, what is AI for business? Like I said, how we think about AI for business is a little different. For a lot of our peers, it is front office, middle office, back office transformation, and we do some of that. But we are picking two lanes. Hey, when OpenAI started, they said we'll solve world hunger problems, but they're now doing more mundane things. But that's solving world hunger is still an opportunity. How do we solve big problems which are unsolvable otherwise without AI? And solving big problems can give big returns. So that's the first focus on AI for business. The second focus is entering completely new markets. Now it could be like new segments, but the big picture is this. There are some services where the product that somebody offers can be bent so much by technology. Either the cost comes to zero because of AI, or the product itself can change so much because of AI. In those businesses, if we otherwise have a right to win, if you have the domain, we have clients, we have access, we want to think of being in that product, in that layer of business, rather than helping our customers transform the product. And actually, we're going to talk about one such later today

which is the market data product. You know, tokenomics is determining IT budget. So let's go back to budget. Top item, human labour. Next, SaaS. Third, cloud and hardware. Everything else is immaterial. But now there's a fourth item which is going to become material. Already, yes. That's token cost. Now frankly, I think deflation for our industry equals to token cost. That's the money that goes to token is going to come off labour spent. So, what does that mean for us? It means two things. How can we participate in that line item of revenue? And that is through smart pricing. And I will talk about that. The second is we don't want to be, we want to be on the side of our clients. What it means is if you spend more tokens, it should not mean more revenues for us. We have to find ways of adding value to customers to optimize their token spend. So, what we're doing is we've launched multiple pricing models. Actually, I'm not actually showing the models here because this deck is public. The core concept is this. Historically, our proposals have what you see on the left column, which is human-only cost. All of our proposals increasingly will carry an option for a client to add the token cost. And we will do it in creative ways. Let me give you a couple of examples. One, we will fix price token cost for some outcomes. This in and of itself is big value because many clients have no clue what it'll cost them. The second way, example, we could gain share of tokens. What do you spend? Whatever you would have spent. If there is a way to measure how much you would have spent versus how much you actually spend, let's do a gain share of that. So, but there are other models. There are eight, 10 different models we are trying with pricing. I'm asking the sales team, like every single proposal that goes from now has to have an option, alternate model. Now, why are we able to go with confidence in a model like we will fix price token cost for you? There are three things we do which give us that confidence. The first is that we engineer for optimal token use. So, think of what that means. You know, you are replacing a workflow software with custom built software. And when you do that, you can custom build static code or you can build it in agentic workflow. You do static code, it uses tokens once to build the software. When you build it as agents, it's using tokens in run time. So, what do you choose to do in agent? What do you choose to do in static? And actually, a platform does this automatically. So, engineering for token efficiency. Second, world-class harness. You know, we, RCT publishes many blogs on tokenomics. Medium, LinkedIn, you guys get a chance, you should read it, it's long reads. But the headline is this, that there are so many models. The difference in capabilities between open and closed models is converging. It's maybe only 5% in some cases. But the price difference is 20x. So, it doesn't mean that that 5% difference in capability is not valuable. But it's not valuable for every scenario. So, one of the things, our harness does many things, but one of the things it does is on a real-time basis, figure out what to use for, what model to use for what purposes.

The last thing, we build SLMs for clients. There are many reasons to think about SLMs, but one of them is tokenomics. So, it's a combination of these three that allow us to be confident in coming up with different pricing models and add value to our clients. By this harness, I said pick different models. You know, there are clients who may say, I've standardized on, whatever, OpenAI, I've standardized on Anthropic. Even if a client has standardized on a single model. There's an example for one of the large labs. These are nine levers in a harness for optimizing spend, even if the client has standardized on a single platform. Now, all of this, for us, has to be backed by talent transformation. We can't do this with the same teams. And in the past, you know, talent transformation was focused on the bottom 95% of the company. You know, you train your teams on digital, you're good to go. I think that's the easiest part of the problem. We have to change top-down, starting from me, every single person in the

organization. I think if you start at the bottom, that is still important, everyone. I feel like we are the most AI literate company in the world. The proof point for us is quadricolor conversation. Our people talking, you go to California Taco Bell, people are talking about models. You go to New Jersey; they're talking about football.

So I think the conversation for us has changed as an organization. People are talking about transformers and weights and models. That's the language in the company. But that's only the start point. I think there are new skills. Now, there's a lot of talk of FDE. I think, you know, FDEs can't execute this vision. FDE can maybe lay out some parts of it, so we think there are at least three layers that is going to go into this. The first is FDEs. Second, equally important, is AI architects, and then everyone else. And so we are, our current goal is to get to 250 FDEs and 1,000 AI architects in 12 months. Three days ago, early this week, we graduated our first internal training program for AI architects. Our internal people, 50, a lot more started, 50 graduated, out of which 10 graduated as what we call as distinguished architects. We will double this output. Every few months, we will have 100 AI architects created internally. FDEs is still going to be a mix of hiring and some internal creation. Client teams' expectation, you know, some things with clients haven't changed, right? How do you win business? You have to have trust with your clients. You don't know their domain, you don't know their business, you don't know their culture. But that building trust today, the pathway to do that is AI. You have to be a thought leader in AI. And so, there's extensive retraining, reskilling of our client teams. Not everyone will make it. We'll change where they don't. And finally, my expectation from our leadership team is different. What succeeded yesterday will not succeed. My expectation is we'll operate at 10x velocity. That's the first expectation. My expectation is if you want to be in completely different business lanes, you have to experiment, you have to take risks, you have to find room and time to think and experiment, not just run what you currently do. This is the whole, all the different aspects of talent.

I want to end with what is, I would say, early outcomes of this, which is about 50% of our revenues are right now AI or AI-infused. Actually, a little more than 50% of our revenues are AI or AI-infused. Now, this is very consistent with what we want to do in AI for IT. Every client, every single day should be positively impacted by AI. This is the proof point for that very rapid progress in whatever time since AI came, we're already at another 50%. Once again, thank you all for being here. We have, following this, we're going to have our leadership team walk you through each of the zeros and some of the components. So we're going to start with Sid, who's going to double-click a bit on the Zerovity platform. We're gonna have Vibhav after that talk about zero license. We're going to have Sanjay talk about two zeros, zero backlog and zero defects. Then we're going to have Girish Ravinderan talk about zero tickets. And Girish Pai is going to round off with one of the components of Infinite Trust, which is data. And then we're going to have, we will break for lunch at some point of time in this. That actual agenda is there in front of you. We're going to have four of our vertical leaders, Shantanu, Ravi, Kush and Ravi talk to you about AI and business.

And we're going to have Jit, our CTO, talk about innovation. These are the AI sessions. We have one non-AI session. I think, you know, to outrun the deflation, one important ticket is improved hunting. And Shantanu is going to talk about; he just took over that new role. We also have new leaders in Europe and Asia. Actually, the Asia leader is not speaking, but he's in the room. If you want, you can meet him, Vijay. Shantanu's going to, however, talk about

how we are rejigging hunting once again. And we're going to finish with a fireside chat, Vikash, VC and I for a few minutes. So, that's what is ahead of you. Sid.