

Eravi-Gopan-President-and-Global-Head

Technology-Products-and-Platforms-AI-day-

Session-8

It's always difficult to lead after lunch, but Shantanu hopefully has woken you all up. I'm Eravi. I'm a little younger than Shantanu in the company. I've spent around nine months in the company. I live in a place called Cupertino, Northern California. I'm the President for Technology, Products, and Platform, which is a new business unit that Hexaware is incubating. My career has been between big tech, startups, and IT services. Seven, eight years before Hexaware, the last seven, eight years, I was first the CRO for Google Venglers product company, and then I was the first CRO for Axel Venglers AI startup before I took up this role in Hexaware. And before that, I spent almost 19 years in Wipro. I was part of the team that built out the technology, media, and telecom business for Wipro, which is a substantial part of the revenues right now.

Technology, products, and platform is, you know, Hexaware meeting this moment in time in terms of our effort to build offerings and a book of business with big tech, hyperscalers, B2B platforms, B2C platforms, as well as software products, SaaS companies, gaming companies, industrial tech, as well as consumer tech, right? That's kind of like the vision and ambition that we have here, making a substantial business unit in the next three to five years. And that brings us to the topic of the AI for the business of AI, right? If anybody is making money at this point in time, it's folks that are in the business of AI, right? Everything else is a promise of the future, and we see that as a big opportunity for us, right? But then, it's a net new area for Hexaware. These are not things pre-2026 we even thought about, but it's a very large economy to go after, and it's a massively growing economy to go after, right? Almost on a weekly basis it's growing, right? Not even on a monthly or a quarterly basis. We sharply focused on two areas. The first is Foundry Labs and consumer electronics, right? Which is, Foundry Labs is a mad race which is going on right now between American, Chinese, and probably the rest of the world is going to catch up on the action pretty soon.

So, you have a fairly large ecosystem of Foundry Labs, model creators, and then on top of that, AI wearables, AI devices, products, which actually deliver specific use cases and services to either consumers or enterprises. There, the focus is really making this work for all of humanity, making it work for all kinds of enterprise use cases. Some of it Shantanu spoke about, right? That's a mad dash. It's a dash against time, dash against competitors, and economics matters, because you heard Keech talk about tokenomics, so doing all this efficiently, not just throwing more and more capital behind it is not going to solve the problem, right? I think Chinese model makers are refactoring the whole industry in terms of what economics means. That's a big focus area for us. It's a picks and shovels opportunity.

The second focus area is AI infra providers, right? The semiconductor value chain, the fab providers, data centre providers, high-tech manufacturing companies. That's a big opportunity. There is an interesting opportunity, because for them to, it's a once-in-a-lifetime opportunity for them to grab the market share when it comes to memory or compute and storage and all of that, but however, you can't create that physical capacity overnight. It takes years to do it, right? Expertise is very difficult to come by. If you look at semiconductor fab,

those kind of niche areas, it will take many years to build and grow the expertise required to do this kind of work, so the only way that they can actually move fast and grab the market share is to use AI to unlock capacity, unlock yield, create more productivity. The interesting thing about this is, one is it falls into that framework that Shantanu mentioned, which is it's net new market, and on the right side, we are trying to solve a hard problem, right? Not an easy problem, using, creating an AI workforce.

The other thing is that it's affecting revenues. It's not about cost; it's not about efficiency. It's both about grabbing more market share and being relevant in the marketplace.

Let's dive into the first one. So, this is the offering that we created at the beginning of the year. Data annotation service has been around for a long time, right, because seven, eight years back, the ML, the machine learning journey needed data annotation, and then, of course, the last few years has only accelerated it. But we see a big opportunity after the model has been created, the post-training part of it. Making it work in as many contexts as possible, using humans to evaluate those models and give feedback and perfect the models. Red teaming to break the models, look at the edge cases and see what is the boundary conditions for these models and use cases, if you must, right? We see a big opportunity in that space, and by the way, it's a proven opportunity. There's a ground sub TAM of \$27 billion. A lot of new age services companies have come up in this space. Cater, AI, Snorkel, Invisible, there's still, the list can go on and on. These are growing very fast. Some of them are very well funded. Some of them are multi-million dollar revenues. However, where Hexaware really brings to the table is the fact that we're able to bring process, tech, and domain together very well. We have perfected the global delivery model. We perfected the model of guaranteeing outcomes, measuring the quality of our outcomes and improving the quality of our outcomes. I think it's not just about talent and models. It's about all these capabilities that we innately have that also matters in addition to that. So, we created this offering. We took it to one of our big tech customers, the Foundry Lab side of it, and we won a pilot deal. Then we did that well. We scaled it up now. We are doing RLHF and a few other things. I'll just play a video to tell you the story behind it.

"In the AI adoption journey, the focus has shifted to achieving real results and tangible outcomes. So, the opportunity now is not just building smarter models. It's building AI that can address domain and cultural context through advanced pre-training and high quality post-training. That's the work Hexaware is doing with its customer. So how do you make an AI speak 1,600 languages? How do you take it from a niche view to all of humanity? How do you build the first billion user voice-driven AI? That's the challenge one of the world's largest AI foundries brought to Hexaware. It starts with data. Millions of lines spoken, translated, and back again. Somewhere in there, meaning gets lost. So, before anything ships, it goes through a pod. Technology, people, and process working as one, led by a principal linguist, a PhD, who doesn't just know the language. He is the language. It's grammar. It's rhythm.

It's grandmotherisms. Every response gets scored through a structured, purpose-built, technology-driven process. A greeting that reads formal when it should feel warm gets flagged, corrected, and scored again. Wrong grammar, wrong tone, wrong assumption. Because incorrect isn't useful on its own, why it's incorrect is what makes the model smarter. That's reinforcement learning with human feedback. The same rigor shapes voice, data, meaning, safety. Because getting language right isn't about words, it's about how the words

are heard correctly. And with care. This is Hexaware. Helping the world's AI finally learn to listen, not just to language, to people. Here is to the first billion voice users of AI. We're grateful to play our part alongside our customer."

This customer has an estimated \$1 billion of budget for the next 18 months doing this and similar work, right? And all of this was taken by boutique vendors for the first kind of global provider to crack into the space with them. They've been pretty happy with the way they're progressing, they requested us to consider setting up a dev centre for them in the US to do this kind of work because they want to collaborate in the time zone. So very interesting models. And it's rinse and repeat. We can take it to every other foundry lab, we can take it to consumer electronics companies trying to launch AI devices and AI wearables.

Come to the second opportunity, right? The AI infrastructure provided. Now that's a broad space. So the beginning of the year we decided to place a bet on high-tech manufacturing semiconductor space. We hired people, we partnered with people, we built the expertise, and we actually opened a Fortune 500 logo, which is a very fast-growing company in the memory semiconductor space. The CEO of this company has come out with an AI workforce strategy, right? So pretty much their stock has gone up significantly year on year, big order backlog, and this Hexaware is a product owner for this AI workforce strategy, right? Three important pillars. One is Euro meetings, which is where their deep experts in semiconductor and fabs actually spend quite a bit of time analysing data, reasoning on data, making decisions on design, on manufacturing, how to unlock that type. Can agents and AI agents be used to actually reason on data and make decisions and elevate the human beings as and when required? Second is, this had to be an enterprise-wide initiative, very grounds up. How can you enable every employee to create AI agents and AI workforce? And that's embedded scout program, which has got a framework to it created by us and also technology competence, which enabled it to be run in a compliant fashion. And then, of course, zero tickets, because massive manufacturing operation, and even otherwise a big supply chain operation, tech operation, whatnot, they're driving the zero ticket initiative.

This is a longish video, it's around three minutes, but it's very rich. It will give you a real-life example of how we are executing the Zero VT vision. So bear with us, it's a three-minute video, but it really shows you the details.

"Let me tell you about one of the biggest waves in technology, the AI CapEx super cycle. Our client, a Fortune 500 high-tech major, is riding it, stock up over 300% in a year with a multi-year backlog from AI infrastructure providers. And they chose Hexaware with our agentic platforms, Agentverse and TensiOps, to lead an AI-driven business transformation.

But step onto their factory floor, and you feel the strain. A global, multi-site, multi-disciplined operation, a massive homegrown platform, 300 applications running the lifecycle. A tool goes down at 2 a.m. The one expert who can fix it is asleep, three time zones away.

So, a ticket gets raised and everyone waits. Backlogs, handoffs, holidays, shift gaps. And the business feels every minute, lower yield, longer downtime, lost throughput, and shipments that slip. The human bottleneck has a very real cost. The transformation begins not with technology, but with value. Every candidate use case is scored across four dimensions, dollar impact, technical feasibility, implementation risk, and time to value. That discipline surfaced the first 100 high-value use cases, each one a business problem with a dollar sign attached. With the targets defined, we advanced on two speeds, an enterprise AI workforce strategy, a harness that industrializes agent engineering with proven frameworks, repeatable playbooks,

and AI-led delivery. Agent code generated, not written, evaluated against hundreds of test cases, and certified through attestation before production. And in parallel, a zero-ticket model, autonomous agents resolving factory problems the moment they appear. And here's what that means for the people. No one is replaced, everyone is elevated.

IT managers move from firefighting tickets to orchestrating agents. Developers from bug triage to architecture, ops engineers from manual monitoring to system design, and factory leads from reacting to predicting. Now, see the platform in operation. In production, the system proactively surfaces everything that matters, connecting the dots from the factory floor to the yield outcome. One real-time command centre, jobs orchestrated end-to-end, human judgment reserved for the edge cases, full traceability, and a knowledge base that compounds with every job, resolution time from 11 hours down to seven minutes. And the returns are already measurable. Experts redeployed to higher-value work, yield protected, shipments accelerated. With phase one alone, two use cases already rolled out, we estimated \$12.5 million a year, and this is only the foundation. As phases two, three, and four roll out, 57 use cases across nine factories, the value scales into the hundreds of millions. So, what does all of this add up to? A factory floor that fixes itself, experts freed to engineer the future, and a business that manufactures more, captures more market share, and grows revenue without a single dollar of additional capital expenditure. In the AI CapEx super cycle, intelligence is the new capacity. Start small, scale fast, ride the wave. Hexaware, creating smiles through great people and technology."

Customer has scaled up beautifully in the last five months since you started working with them. The Zerovity vision, we created this, the business value framework, the workforce strategy, inspired by our blueprints in Tensai Ops and AgentWorks, which is part of Zerovity core. And now they started using RapidX for app modernization, Amaze from tech modernization. So, this really is a customer having a journey with us, using our full suite of capabilities, and hence being a real-life example of how zero-VT could be quite powerful for our customers. Bringing it all back to TPP, the business that we are trying to incubate, this is the year we are building the foundation, building the offering, the core book of business and the teams. We feel good where we are. We closed around 25 contracts, YTD. At least eight of them are purely the AI for the business of AI kind of space. Of course, many of them start small and then scale. We're doing a really good job of mining our existing accounts. We won a fairly important framework engineering outsourcing deal with a big tech company, which is not something typically we would have done in the past. We are in the final mile, fingers crossed, for being the gen-AI services provider for a big tech company, right? So really tangible progress. The pipeline is progressing really well. Where we want to head, I think we feel really good about it. Probably closing around two to three deals a month, right? So good, solid progress. Of course, the Bay Area centricity is there to the business, building a good core team there, delivery capabilities in Bay Area and the US, but also, it's truly global because these companies are now trying to address the world.

So, thank you very much. And I would like to invite my friend and colleague, Ravi, to talk about AI in financial services.