

Girish-Pai-Executive-Vice-President-and-Global-Head-Data-AI-day-Session-6

Trust, which is one of our foundational pillars that helps drive zero friction, I want to tell you a little bit about what we do in the data space. We offer the full spectrum of services that our clients require to run their data, starting with strategy and advisory, where we help clients figure out what their data strategy should be. We have a very large data engineering business, where we help clients build both greenfield data estate, as well as help them modernize their data estates that they have accumulated over years. We also offer data management services, which includes things like master data management, data governance, lineage observability, which is emerging as a key capability in the world of AI. And lastly, we build business solutions, where we help clients solve business problems through things like data science, analytics, and insights, where they leverage the power of their data to drive results. But today, I will focus largely on what is AI, how is AI impacting the data space? And we have three focus areas that we work on.

The first is data readiness for AI. And this, as Keech mentioned, is a key pillar of building infinite trust. Enterprises have come to realize that one of their key differentiators in driving value out of AI is their data. The results that they can achieve is directly related to A, how good and reliable their data is, and B, how effective are they in taking that data and making it usable and understandable for AI solutions. For the rest of the presentation, I will largely focus on this area. But before that, let me give you an overview of what the other two are.

On the left of the slide, you'll see what we're simply calling agentic data engineering. And what we've done here is, as I mentioned, data engineering is probably our largest subservice offering within data. And what we've done here is we've taken what we have traditionally done across what I call the data supply chain, where data is originated in source systems all the way to where it is consumed by end business users. Either in the form of dashboards, or reports, or charts, or more sophisticated business users who are taking raw data and performing analytics and data science on it. But every step of the way, what we've done is we've taken what our consultants do, whether they're engineers, they're ingesting, they're data modelers, they're data profilers, or they're designing business insights reports. Every step of the way, we've taken what they do and modeled it and seen, how can we use agents to A, amplify their capability, and or B, bring more efficiency? Today, when a client hires a data engineering team from Hexaware, our consultants don't show up with just a laptop.

They also show up with an AI agentic assistant. And the beauty is that we've built these assistants so that they are compatible with the client's environment, no matter what it is. So, it could be an agent that we deploy, or we have a pattern or an accelerator readily available that we can reassemble the agent in the client's environment using the technology stack that they have.

Through this, we are seeing somewhere between 40% to 50% efficiency improvement, and 25% to 30% commercial benefits being delivered to clients. Now, most of our clients have a situation in the data space where their backlog of work significantly exceeds the capacity that they have allocated. And so, our clients are able to use our agentic data engineering features to knock things off their data backlog in a more and more effective way. The last area of focus is on AI solutions. And there's a term that we coined called vibe analytics. Gartner calls it

agentic analytics. And it's a key focus area that they're expecting significant growth in. What we believe is that the way in which business users, especially executives, consume insights is changing. With the ability of AI to rapidly do things, to rapidly uncover insights, Sanjay talked about how AI is eating the software world. AI is also starting to impact the data science world. So, in the old days, when you would go to a data science team with a question and expect to hear insights in weeks, we're able to do that in real time through a conversational interactive experience that we're simply calling vibe analytics. Think of the problem you want to solve and describe it. And vibe analytics will use your data to solve it for you. It'll tell you why you're seeing what you're seeing. It'll tell you what you might do about it.

For the rest of the presentation, I will talk about data readiness. As you must be aware, almost all enterprises are engaged in POCs with AI. There is a mortality rate in terms of even successful POCs not making it to production. And what we've seen is 60% of that mortality can be because of causes that are related to data. And it's not that they didn't have data in the first place. It's just that the effectiveness of them being able to use the data for AI solutions is limited. They're not able to drive the promise of the value that they expected. Expectedly, clients are starting to spend money in terms of modernizing their data estate, in terms of making sure that the data that they have are usable for AI solutions. The estimate for the size of this industry is \$40 billion. But the more exciting part is that this is expected to triple over the next five years. So very, very promising. Now, it's not an easy problem to solve. To make your data understandable for AI is not simple. And I'm going to call out, this slide has six different aspects, but I'm going to call out a few that I think are the most important. The first is that while knowledge resides, facts reside in data, for true knowledge, you need context. And context doesn't sit in tables. It doesn't sit in databases. It sits in documents, it sits in policies, it sits sometimes in the mind of the person that's interacting with the data. And an enterprise's ability to codify that context is very, very important. The second is most clients have a fragmented data environment. It could be based on geographies, it could be based on business units, but their data resides in silos. And the effort required, the engineering and human effort required to coherently build knowledge across those silos is sometimes very significant.

The third area that I want to call out is on the bottom left of the screen. Most data states, data structures, were created with a human being in mind. When we created data warehouses and reports, the expectation was that there's going to be a human at the other side who can read a chart and understand what it means. But that premise is getting dispelled as we start looking at a proliferation of agents. Humans intuitively understand charts, they intuitively understand when they look at a data source, they kind of know where it came from and they intuitively understand. But agents and agentic solutions don't. And so, there is a need for data sets to be purpose-built for consumption by agents.

Now, many enterprises have grappled with this. Many of our clients have gone down the path saying, I'm going to solve this problem. And most clients will throw up a chart, a blueprint of what their future data estate needs to look like, and in there you will find a context map or a contextual layer or a semantic layer. And there's nothing new there. The problem is that in order for a client to actually go down that path, there are many hurdles they need to overcome. The first being there are already so many projects going on and to embark on another one to unify their data state and create a semantic layer is Herculean. There's

completing priorities. The second is it's not a small task. There is no magic wand. It takes anywhere from 12 to 18 to 24 months and significantly expensive to build.

The third, and I'm not going to read every single one there, is there is a shortage of talent and a shortage of skills. And back to the question about completing priorities. The last one, trying to combine the knowledge that resides in both your structured and unstructured data requires a level of coordination that is difficult to master. So a typical client, when they're thinking about something like this, is faced with embarking on a project that's 18 to 24 months and costs tens of millions of dollars. Now we have tried to solve this in a much, much more effective way. Now my colleagues who came up before me talked about what we're built with Zerovity, and you've already heard of some of the features that comes with it. What we have done is we've taken those capabilities and through an agentic scan, AI-led human-assisted agentic discovery of the client's environment where we look at data warehouses, databases, applications, documents, other sources of unstructured data, as well as to some extent knowledge that might be in human minds. We use this to build a contextual map of the client's environment. And this map includes two key pieces. One is the semantic knowledge layer in which we embed context, curated data sets, and the overlay of structured and unstructured data. The second is the model intelligence layer where you can run industry purchase models or your own models, but you're able to integrate into it your own data. The beauty is, and I'll show you a short demo of how this works, beauty is we can do this for one functional area in as less as six to eight weeks.

The second thing that we've done in the spirit of the zero friction is that we're trying to build this in a way that it leverages any investments that the client has already made. For example, they might be down a path of modernizing with Databricks and they may halfway through that path. We're not going to require them to change anything in their stack.

They might have built a semantic layer for one part of their business. We will bolt on to what they already have. We're able to do this in a matter of six to eight weeks.

If you think about that, the choice that they have is go down a multi-year program spending tens of millions of dollars or do this in iterations of six to eight weeks costing a million dollars each. What I'm going to do now is show you a quick teaser video of how this works. During the floor walk, I look forward to meeting many of you in our booth where we will have a longer version of our platform demo-able, but I'm just going to show you this quick teaser video.

"Making data AI ready from discovery to semantic understanding. Every enterprise captures data. Most of it is not AI ready. This is not just a data problem, but also a business problem. AI initiatives stall because they're missing four things, context, coherence, intent, and trust. To deliver real business value enterprises need to address this so AI outputs can be trusted.

Delivering this is a massive multi-year effort for enterprises. Here is how Hexaware turns your data estate into an AI ready foundation. Hexaware's agent powered data estate discovery autonomously scans, analyses, and catalogues all data assets of an enterprise.

Databases, warehouses, BI platforms, pipelines, mainframes, even policy documents. No source is out of reach. Take Viridian Global, a bank shaped by decades of growth and acquisition. Discovery ran across four systems, legacy, cloud, and BI, surfacing definitions that didn't always match. This is Viridian's data estate map, a live view of every KPI measure traced back through its data model, its source table, to the original system. So every number can be verified. That's what builds trust in the data. Where definitions conflicted, only one moved

forward, becoming the single definition every application now draws on. That's coherence. One truth consistently applied. All of this builds the semantic knowledge layer comprising a glossary, taxonomy, ontology, and a knowledge graph. At the centre of the knowledge graph sits the bank itself, expanding into regions, business domains, entities, and metrics down to the lowest level of detail. This gives AI the context and intent it needs so it understands not just what a number is, but what it means. This same graph powers multiple use cases, each grounded in the same resolved truth. Ask Hexaware's Vibe Analytics a real business question, like a lone officer checking exposure to their client, the ABC group.

With a knowledge graph behind it, Vibe Analytics delivers a governed answer instantly, understanding context and intent even when they're never explicitly mentioned. Every enterprise captures data. Now yours has context, coherence, intent, and trust. An AI-ready foundation delivered in a matter of weeks."

I'm going to round off with a couple of examples of how we've deployed this. The first is at a investment pension and insurance provider. For this company, for the investment side of their business, their data was scattered across 10 plus silos. And their executives were frequently trying to model what market events might have as an impact on their investments. And typically, something like this took weeks to get done and required significant manual data preparation before they could run any sort of what-if scenarios on it.

Using our Zerovity tool, what we did was we first built a modern data lake for them on GCP. On that, we implemented a semantic knowledge layer and built a few different data products that were designed for the kinds of questions they had. And by implementing this, what we've realized for them is there is a 70% reduction in the amount of manual effort that's required for them to do some of their what-if scenario modelling. There is an 85% improvement in time to market, especially when they're looking at market events and how they could potentially impact their investments. And if you think about the world today, how many market events that are happening, they're able to respond to those significantly faster.

The second example is for a different client. And before I get there, while the primary purpose of us designing the solution was to create that enterprise map, that contextual understanding or contextual intelligence of the enterprise, what we realized is that knowledge actually can help power other use cases. And one of those other use cases is the modernization of your data estate. And so, the second example is for a big four audit firm where they had literally 5,000-plus tables across countless SQL Server databases that they were running their business on. And what we were able, and previously, while they wanted, they had an ambition to unify all of that in a single enterprise data lake, the cost of it was, and the time it required was prohibited. It was going to take at least two and a half to three years. And by automating the scan of their environment and being able to build the understanding of what they had in place and then applying our intelligence on how we could make that significantly more efficient, we were able to modernize them to a Databricks stack, displacing a bunch of their legacy technologies and standardizing on Databricks with Power BI. We were able to do that in less than a year with somewhere between 40 to 60% efficiency improvement in the migrations. That's all I have for you today. And I believe we are ready to adjourn for lunch. So thank you for your time. Welcome once again to Hexaware.