

Siddharth-Dhar-President-and-Global-Head-AI-day-Session-2

Good morning. My name is Siddharth Dhar. Everyone calls me Sid. And I head AI for Hexaware. So, you heard Keech talk about zero friction being powered by Zerovity. You saw that in the video as well, right? And please trust me, you'll hear about Zerovity at least 100 more times the rest of the day, because our propositions are powered by it. So, I thought it's great to spend a few minutes upfront to sort of explain what Zerovity is. What does it do? How is it different from others? And how does it help us? So, if you think about it, one of the most open secrets in enterprise IT is that most enterprises have very fragmented data and knowledge of the applications that run their business. The reason for that is a typical Fortune 500 company can have anywhere from 5 to 5,000 applications. Most of these are built over the last 30 to 40 years. They have undergone thousands, if not tens of thousands, of changes. So to keep that documentation current, to keep that context current is an impossible task. But it is important, because without that, it's very expensive to make change, it's very risky to make change, and it takes very long to make change. So that's the problem we want to solve with Zerovity. In fact, if you look at any of the pillars of zero friction, the root cause for many of those pillars is this problem, the fragmented knowledge of your own application environment. In Zerovity, what we have built is the capability to ingest data about the application environment and then use that to recreate a living map of that environment. We call it the living context.

And we don't do this by simply consuming the fragmented documentation. We actually go to the source itself. We have created parsers that help us parse every single line of code and use that to recreate business context, recreate how applications interact with other applications and the environment around them. And this is really tough. You can't do that if you don't do the hard work. You may be thinking for a second, hey, we have heard this story before. Every Frontier Lab says, my model can comprehend code, my model can reverse engineer it, and my model can forward engineer a new application out of it, which is true in many cases. However, there are three clear differentiation points for Zerovity. The first differentiation point is it's not about being able to comprehend code using an LLM.

It's about how accurately can an LLM comprehend that code. All of you have used LLMs. Right? The one warning that an LLM comes with, right, all of us see that every day, that not everything that an LLM gives you is correct. That's a warning every LLM gives you. So, you can't use that to get a fully 100% view of your business context. You need to do something more, which is actually the code comprehension piece, right? So, you need to build parsers that understand code and use that to build that context back up.

That's a big differentiation. We do not tell, we do not rely on AI to tell us the context. We uncover the context and then give it to AI for AI to do better things with it, right? That's the first point of differentiation. The second point of differentiation is we're not interested in building simply a knowledge map of an application. There are many products that do it. Our real intent is to do that for every application and then understand how each application interacts with the next application and how do they truly deliver the business outcome.

And that is, again, very tough to do, right? You really need to work hard to be able to uncover that level of context. But without that, the context is incomplete, right? So that's the second point of differentiation.

The third point of differentiation is that we have actually built an agentic harness into Zerovity. In fact, it is the same harness that we use in RapidX. It's an improved version of that. And we already know from our customers, and customers have tested our RapidX harness against our competition, against even frontier models, we already know from their evaluations that our harness is more token efficient and delivers better quality of code, right? So that's the third advantage that we bring with Zerovity.

One quick point I want to make here, if you look at the slide in the back, and I think Keech alluded to that as well, Zerovity will replace all of our existing platforms, right? So, Amaze, RapidX, and Tensai will all be replaced by Zerovity. We have made it backwards comparable, but I think over the next few quarters, we will go back and replace all of those platforms in every customer where they're installed with Zerovity. That should give us a jumpstart in the installed base as well.

The final thing I want to talk to you about is how do we expect customers to buy Zerovity? Actually, we don't expect customers to buy Zerovity. We expect customers to come to us with one of the points of friction that Keech talked about. Some customers will come back and say, hey, we have SaaS sprawl and we want to get rid of it. Some customers might want to say we want to modernize. Some customers might want to say that, hey, we cannot cope up with vulnerabilities. No matter what pattern you come to us with, our answer begins and ends with Zerovity. So if we answer your friction for one point with Zerovity, and Zerovity is already established, it then allows us to go and talk to you for all other points as well. And that's what creates infinite momentum. And to me, that is the true superpower of Zerovity.

With that short introduction to Zerovity, I'm going to hand it over to my colleague, VB. He's going to talk to you about our current favourite topic, Zero License.