



E2E Networks Limited

CIN- L72900DL2009PLC341980

Uppal's Genesis A-32, Block B, Mohan Cooperative
Industrial Estate, Badarpur, New Delhi 110044

Email: cs@e2enetworks.com, **Website** <https://www.e2enetworks.com/>

Date: April 22, 2026

Listing & Compliance Department
The National Stock Exchange of India Limited
Exchange Plaza, C-1, Block G,
Bandra Kurla Complex, Bandra (E),
Mumbai, Maharashtra-400051

Ref: NSE Symbol- E2E

Sub: Transcript of Analysts/Investor Earnings Conference Call for Q4/FY26

Dear Sir/Madam,

Pursuant to Regulation 30 of Securities and Exchange Board of India (Listing Obligations and Disclosure Requirements) Regulations, 2015 and other applicable Regulations, please find enclosed herewith the Transcript of Analysts/Investor Earnings Conference Call, which was held on Monday, April 20, 2026, at 04:30 PM (IST) to discuss the Company's Q4FY26 earnings.

The aforesaid information shall also be made available on the website of the Company at <https://www.e2enetworks.com/>.

Kindly take this on record.

Thanking You,


Yours faithfully,

For E2E Networks Limited

Ronit
Company secretary & Compliance Officer
ICSI M. No.: A59215



“E2E Networks Limited
Q4 & FY26 Earnings Conference Call”
April 20, 2026



**MANAGEMENT: MR. TARUN DUA – MANAGING DIRECTOR – E2E
NETWORKS LIMITED
MR. NITIN JAIN – CHIEF FINANCIAL OFFICER – E2E
NETWORKS LIMITED**

MODERATOR: MS. RASHI KHATRI – GO INDIA ADVISORS

Moderator: Ladies and gentlemen, good day, and welcome to E2E Networks Limited Q4 and FY26 Earnings Conference Call. As a reminder, all participant lines will be in the listen-only mode and there will be an opportunity for you to ask questions after the presentation concludes. Should you need assistance during the conference call, please signal an operator by pressing star then zero on your touchtone phone. Please note that this conference is being recorded.

I now hand the conference over to Ms. Rashi Khatri from Go India Advisors. Thank you and over to you Rashi.

Rashi Khatri: Thank you and good afternoon everyone and welcome to the Q4 and FY26 earnings call of E2E Networks Limited. We have on the call Mr. Tarun Dua, Managing Director of E2E Networks, and Mr. Nitin Jain, CFO of E2E Networks. We must remind you that the discussion in today's call may include certain forward-looking statements and must be therefore viewed in conjunction with the risks that the company may face.

I will now request Mr. Tarun to take us through the financial and business updates, subsequent to which we can open the floor for Q&A. Thank you and over to you, sir.

Tarun Dua: Hi everyone. Hi to all our team members, all our investors, and we hope that you have had a good last financial year. I will do a brief coverage of all the progress we have made during the last quarter and last year, and then I will hand it over to Nitin to talk about the financial highlights of this quarter and the overall year, and then we will open the floor for questions. In this year, we have had a lot of milestones as a company and as a team that we have achieved.

We have been able to use our TIR platform for large-scale GPU clusters, and we have helped LLM teams operate trainings on large clusters over the TIR platform. With this, we have demonstrated a full stack capability of operating both bare-metal and container-based large-scale deployments for NVIDIA-powered GPU infrastructure at a bigger scale and we have like not only demonstrated that like we have the ability to operate the infrastructure but also successfully monetize the infrastructure at a much larger scale than we had proven like last year.

This is a testament to the strengthening of our technology and talent, which has been used to build very deep in-house capabilities across the entire five layers of the stack. We have done major software improvements to improve the reliability, performance, and scalability of our GPU infrastructure in the last many quarters.

Going forward, we feel that we have been saying this a lot, that the future is all AI. This is a decadal story. Now, given how the markets all across the world are reacting to the AI infrastructure usage being increased very, very drastically, I think we have been validated in putting up speculative GPU infrastructure. Where the world is divided between haves and have-nots, we are the haves in terms of having access to speculative GPU capacity where capacity is running out everywhere.

That allows us to continue our growth. With regard to where we see medium-term and long-term kind of growing in terms of the world shifting to GPUs from the CPUs, from rule-driven software to AI-driven software, we continue to explore how we can build more infrastructure under our management. We have explored equity, we have explored debt. Now we are also

exploring a variety of private credit asset-light models to bring in rapid expansion to our GPU capabilities.

We have consistently met the industry benchmarks for performance for inference or training or GPU deployments under our management. We expect our Cluster 1 B200 1024 to go live somewhere in the mid-May, and in a couple of months, we are expecting to be able to deploy another cluster of 1024 which has already been planned. Then we are very strongly positioned to capture significant uptake in the Blackwell generation of GPUs apart from continuing to run very, very strong on the Hopper generation.

Apart from this, we continue to build a plan around B300, GB300, and Vera Rubin deployments. We will continue to keep increasing the GPUs under our management and under our direct deployments, both through partnerships, through direct acquisition, and by exploring various financing models that can help us procure more GPUs. We continue to focus on that.

We continue to focus on our end-to-end project management capabilities to deliver AI infrastructure to AI services and AI product companies, AI digital natives, enterprises in the BFSI and other sectors. We will continue to work towards that. So okay, so I think like that pretty much covers it. This has been a good year and a good quarter.

Pretty much our strategy has been validated over last many quarters where we significantly continue to add to our technology heft in terms of having deep-tech talent in-house to be able to do stuff. We gave a certain guidance in terms of where we would be by the end of this quarter a couple of quarters back. I think we have more or less met and exceeded that target, and we hope to continue to do better in the future.

Given that the market is growing very, very positively and it is a broad trend, there can be again, as we have mentioned in the past, that there are crests and troughs, but ultimately this is the market where the GPU utilization is going up, the demand for tokens is through the roof. India being a country of more than a billion people on phones, and of course even if they were not on the phones, there is an inherent ability to treat every possible piece of information that is going in and out of people, in and out of machines, in and out of enterprise systems to be described in the form of tokens.

Tokens can be representing voice, text, any other form of data, whether invisible spectrum, audible voice, inaudible sound waves, ultrasound or whatever. Everything is tokenized data for AI, and I think India as a country generates a lot of data. We continue to believe very, very strongly in the long-term India story of AI factories being built here and India becoming the AI factory of the world.

Now with that, I would like to hand over for more numbers-driven financial highlights talk to Nitin, and then we will open it up for question and answer. Nitin, over to you.

Nitin Jain:

Thank you, Tarun. Good evening, everyone. Thank you for joining us today. I will walk you through our financial performance for the fourth quarter and full year FY26. Q4 FY26 again witnessed a good quarter for E2E Networks. We have demonstrated our ability to scale AI infrastructure rapidly, achieve high utilization across GPU clusters, and convert capacity into

strong revenue growth. Most importantly, we have been able to drive our EBITDA margin in Q4.

Let me start with the quarterly highlights for the Q4. Revenue stood at INR956 million, which is up 186% year-on-year and 37% quarter-on-quarter. EBITDA stood at INR581 million, with the EBITDA margin expanding up to 60.7%. Profit before tax turned positive at INR86 million compared to a loss of INR75 million in Q3. PAT stands at INR64 million. This performance reflects INR120 million swing in profitability sequentially driven by operating leverage and a strong execution.

EBITDA margin expanded by 413 basis points sequentially. At the same time, depreciation increased to INR513 million in Q4, reflecting ongoing infrastructure investments. Despite this, we achieved positive EBIT and positive PBT. For the full year FY 2026, revenue stands at INR2,456 million, which is up 50% year-on-year. EBITDA crossed INR1,263 million, which is up 30.6%. However, we reported a PAT loss of INR156 million, which is driven entirely by the depreciation on our GPU infrastructure investment.

Our core business remains strongly cash positive, operationally profitable at EBITDA level. As utilization continues to ramp up, revenue will progressively outpace depreciation, improving reported profitability. To conclude, Q4 demonstrates our infrastructure investments are translating into revenue and we are firmly on a path for growth. Thank you for your time. Now I would open the floor for the questions.

Moderator: Thank you, sir. Ladies and gentlemen, we will now begin the question-and-answer session. Our first question comes from the line of Bhavya Gandhi from Bajaj Alternate Investment Management Limited. Please go ahead.

Bhavya Gandhi: Yes, hi. Thanks for the opportunity. First question is regarding the asset-light model that we are looking out for with L&T. Could you throw some light, how will the asset-light partnership fructify, what sort of arrangement we have?

Tarun Dua: First of all, of course, we have a MOU with L&T to monetize the GPU infrastructure that they are building. That is still in exploratory stage. As the GPUs get deployed and we start the work into monetization, as we know more and more, we will talk more and more about that. That being said, this is not an exclusive arrangement. We will continue to operate at an arm's length with L&T and we will continue to explore other partnerships of similar or different nature as well.

There is a long gap between pure debt versus pure equity. There are a number of structured possibilities in terms of figuring out how to finance the GPU in partnership with a lot of different types of partners with differing objectives. We will continue to explore all those partnerships. As and when we reach a conclusion on any particular partnership, we will obviously keep everyone informed.

Bhavya Gandhi: Sure, but will it be margin accretive because the capex is going to be on other partners?

- Tarun Dua:** It is too early to say, Bhavya, how would the numbers look like. Obviously, it is very, very early to say how those things would look like. Unless the structure is finalized -- see, obviously we all work for a profit, so we will not do anything which does not earn us the profit. That being said, how to put numbers to that is not something we can do today.
- Bhavya Gandhi:** Got it, sir. And on the MRR, would you like to guide for the next year because you have been doing this capex from your own pocket also this year? If you can throw some light.
- Tarun Dua:** Yes, so we will continue to focus on growth just the same as rest of the market is growing. We want to be very, very growth-focused. That being said, I think it would not do justice to do a MRR guidance for one year in the future. The future is changing very, very rapidly week on week. It was obviously impossible to give a week-on-week guidance on MRR and it would be better to watch it quarter-on-quarter than predicting four quarters in advance.
- Bhavya Gandhi:** If you can just throw some light on the asset turn for the capex that we are doing, at least some sort of understanding we can get in terms of capex.
- Tarun Dua:** We want to get the framing away from things like asset turns. Essentially, don't look at us as an asset monetization business. Look at us as a technology business. I think there was a recent conversation around the tokens becoming more valuable. The same set of tokens that were being generated by say open source AI or closed source AI, what used to produce say X amount of value, we are increasingly seeing that as the accuracy, efficiency, and capabilities of AI increases, then the generated tokens also become more valuable for businesses who have figured out how to utilize those tokens in their business.
- Which means that setting up ourselves for looking at the value of our outcomes based on the size of our assets doesn't do justice to the business. Ultimately, we want to get away from saying that, okay, this infrastructure produces only X amount of value. We will continue to explore options for how to produce more value from the same infrastructure by figuring out where the tokens are more valuable.
- Bhavya Gandhi:** Got it. Fair enough. Okay, sir. That is it from my end. I will get back in the queue. Thank you so much.
- Tarun Dua:** Thanks, Bhavya.
- Moderator:** Thank you. Our next question comes from the line of Keshav from Niveshaay. Please go ahead.
- Keshav:** Yes. Hi. Thanks for the opportunity and congrats on the great set of numbers.
- Tarun Dua:** Thank you, Keshav.
- Keshav:** So sir, we procured the Blackwells way back in the early quarter and they are still not live. What is the reason behind the delay? Are we facing some kind of procurement issues or something from NVIDIA?
- Tarun Dua:** Global supply chains have been impacted somewhat. The delays are always for want of a horseshoe nail. Kind of we have been working very diligently towards making sure that

everything is planned out and everything gets delivered on time, but you cannot control every single component. Sometimes some components can be the most critical ones from the point of view of getting the entire deployment done. We are targeting the first deployment to go live before mid of May. Keeping our fingers crossed over there.

Keshav: Got it, sir. And this quarter, GPU demand has been strong globally and people are saying that all the GPUs are sold out, including the old ones. How are we seeing this trend in the domestic market? Are we also experiencing the same or is it a different situation in India?

Tarun Dua: I do not think I have understood your question, but let me try to answer what I have understood. See, obviously it has turned into a somewhat much better market than it was a couple of quarters back, especially for people who have the GPU infrastructure in place. That being said, as a company, as an Indian company, we are very, very focused on India first. Wherever we get an opportunity to support Indian companies, we definitely like to do that. Now, with that in mind, basically we always prioritize India first and then we are also happy to support the global infrastructure needs today. I hope the answer is in the right direction towards what you were asking.

Keshav: Yes, definitely. I was just trying to ask if demand has been strong globally that the GPUs got sold out. I was just trying to understand that.

Tarun Dua: Yes, yes. There is fresh, we are seeing demand both in India as well as globally. It is strong both ways. We definitely try to prioritize India. But then wherever the demand comes from eventually, we are happy to fulfil it.

Keshav: Got it. And one last question from my side that employee cost increased this quarter, so if you could give some color on that, what is the key drivers behind this rise?

Tarun Dua: See as we grow, we kind of start figuring out more and more interesting problems that need to be solved in this space. Obviously, our goal is to go towards what you call higher value tokens and that obviously requires the application of quite a high level of talent. As we obviously grow, the base effect would still be there. But on the other hand, we want to maintain a balance of not losing out on future opportunities because we did not invest today on what was required 12 months later or 18 months later or two years later.

Keshav: Got it. And one last thing if you could also provide the GPU utilization measure rate if possible, for this quarter.

Tarun Dua: Broadly across our entire infrastructure, not just GPUs, I think we are looking at like 80% plus utilization. Certainly across the quarter, I would say only in the March month, we are looking at 80% plus kind of utilization. Certainly less than 85%. So there is elasticity in terms of basically how we can increase the utilization further over there.

Keshav: Got it, sir. I will get back in the queue. Thank you so much and congrats again for the great set of numbers.

Moderator: Thank you. Our next question comes from the line of Bharat Gulati from Dalal and Broacha. Please go ahead.

Bharat Gulati: Yes. Hi. Thanks for the opportunity. Hi, sir. Just wanted to get a breakup of what is our MRR breakup currently is. Can you give a split between India AI Mission, enterprise, and also a split between what would the GPU contribution of that MRR would be as compared to CPU?

Tarun Dua: I think broadly the way things are progressing, in another couple of quarters, I think GPU contribution would be closer to 85%-90% over coming quarters. That being said, all these numbers are fairly dynamic. In the sense that they vary very, very rapidly from a week to another week. I would not like to and the sample size today is fairly small. So what the split is something that we do not want to worry about today.

As we grow the cluster sizes to a much larger number, then it would start making sense to talk about, okay, Hopper you have got a couple of thousand of them, then Blackwells you have got a couple of thousand of them. That is the point at which it would start making sense to say, okay, what is happening on Hoppers, what is happening on Blackwells. Today the entire universe of GPUs, CPUs, storage, all of that stuff that we have is fairly, fairly small to put a very, very strong split around those numbers today.

Bharat Gulati: Fair enough, sir. But could you just give a MRR split between the India AI Mission and our enterprise or our SME clients? You know, just trying to get an idea of what kind of visibility do we have in that INR374 million MRR? And also just to add on to that, before that our MRR for the previous two months if we average it would come to a INR290 million kind of MRR. So just trying to understand that spike that came in the last month. What was the reason for that spike and also how sustainable is this INR374 million MRR?

Tarun Dua: It is broadly an increase in overall utilization that has led to the increase in the spike in the March month. From the perspective of overall split, I think the India AI kind of -- Yes, so basically the overall government business across this quarter has not exceeded more than say 35%-40%.

Bharat Gulati: All right. So sir, could you is there some particular reason for that because ideally our India AI missions were supposed to start in the month of Jan. We are in March as of this quarter's reported numbers. So what kind of ramp-up are we seeing in that? And just any kind of timelines and what kind of revenues in terms of MRRs do we expect from that?

Tarun Dua: The way we have always looked at India AI mission is that like, look we are trying to obviously like we said that we are supporting India, we are supporting India AI mission, we are supporting Indian companies and we definitely give a preference to all those workloads. But that being said, whatever capacity is not off taken, so we are not trying to force anyone to offtake any capacity from us. We are very, very happy to sell it to outside the scope of either the AI Mission or the government workloads.

We have enough customers across multiple segments to solely rely on one revenue driver growth. But that being said, we continue to work very closely with the AI Mission for the coming generations of GPUs as well. And we are hopeful that we will continue to collaborate, continue to work together, continue to solve problems.

Bharat Gulati: So sir, just trying to understand that India AI Mission was almost all our Hopper series GPUs given to India AI Mission. We are seeing 80% utilization but India is not...?

Tarun Dua: No, we have not scaled up that to that extent. So majorly whatever was the scale required by the allottees of AI Mission, we have provided that scale. But we have not insisted on them having to increase their scale without having the need for that. So we are quite okay with that.

Bharat Gulati: So in terms of our revenue currently, what kind of clients are we currently catering to and what is the pipeline? You know, you said mid-May we expect Blackwell to get deployed. Do we have firm orders?

Tarun Dua: Again, this is a very, very small base. So, it's like, okay, there are GPU customers everywhere in the world. There are GPU customers in India, there are GPU customers outside India. Every possible segment that you can think of, whether it is enterprise, whether it is BFSI, whether it is education, everyone needs GPUs. So, it's practically like -- basically like almost first come first served today.

So like the nature and profile and putting a stat on, okay, this much of our business comes from here, this much of our business comes from there, is very, very premature. I think we need to grow the base of GPUs to a couple of tens of thousands before these numbers start to make sense. Today at a very, very small base, none of these numbers would make sense to track and then kind of say that, okay, what's the kind of numbers this quarter? It will change next quarter.

Bharat Gulati: So sir, would it be fair to say that until we do not build a huge base of GPUs, our MRRs will be lumpy in nature? Is that what you're trying to iterate at?

Tarun Dua: No, I think the demand has secularly changed in the GPU world. From the era in which the lumpiness was there because our base was even smaller, I think we have come a long way over there. That being said, the lumpiness will continue to decrease as we continue to build more and more volume of GPUs under our management.

Bharat Gulati: Sir, out of this INR374 million, what exactly would be long-term MRRs that at least we are seeing a visibility for the next 8 to 10 months or a year or so?

Tarun Dua: See we are not very, very focused on -- so we have a balance of basically thinking about, okay, what needs to be 6 months, 1 year, month-on-month, which is preemptible. So we are doing a whole series of them. So it's a week-on-week effort to figure out that, okay, what is the capacity you want to sell hourly, what is the capacity you want to sell weekly preemptible, what is the capacity you want to sell monthly, what is the capacity you want to sell yearly.

That being said, wherever larger clusters are involved, we try to do at least 6 months to 1 year visibility with our customers, and increasingly for even larger customers, we are trying to look at even longer visibility to at least 2 years to even going up to 3 years.

But that being said, our view is that as tokens become more expensive to generate and kind of generate more value for people. It is better to not put all the eggs in the long-term basket but

have an increase judicious mix in very short-term contracts and medium-term contracts and some degree of long-term contracts.

Bharat Gulati: Got it, sir. Got it. And just one last thing. In terms of the Blackwell that we expect to get deployed in May, can you give any idea in terms of who will be the customer? Will it be an enterprise, India AI Mission, an SME, and when would that customer's revenue start to come in? Would it come from May or will there be a lag?

Tarun Dua: We will announce that after closing that particular deal instead of saying anything speculative today.

Moderator: Our next question comes from the line of Nishant Joshi from Equisense Advisors Private Limited.

Nishant Joshi: Sir, I have one question. As you said company is building up their assets using debt...

Moderator: I am sorry to interrupt you, Mr. Joshi, but your voice is not audible.

Nishant Joshi: Sir, I want to say that as the company is planning to build assets using debt as well as the equity and subsequently intend to go for the asset-light model also. So will this change our business model? Means instead of providing hardware, we would be providing more of service as in we will be offering our TIR platforms. So will this lead to change in our revenue model?

Tarun Dua: Think of this as an expansion of the number of business models without shutting down any of the existing business models. So we will continue to own, continue to operate GPUs, we will continue to acquire new GPUs, we will continue to expand the partnerships to bring increasing number of GPUs under our management available through our TIR stack platform which operates at multiple layers where it is the choice of the customer what layers they want to buy.

Nothing is changing from the point of view that, okay, we are not going to do this or we are not going to do that. We are simply saying that we are going to do additional number of things to expand the universe of our thinking to do a lot more than what we are doing today.

Nishant Joshi: So sir, it would be an ideally additional line of revenue you mean to say when we move?

Tarun Dua: Absolutely. Absolutely.

Moderator: Our next question comes from the line of Deepak Poddar from Sapphire Capital.

Deepak Poddar: So just wanted to understand now at a March MRR of INR37.4 crores. So what is our capacity utilization and on what capacity?

Tarun Dua: So this is closer to overall capacity utilization of around 80% in the March month.

Deepak Poddar: And on a capacity base of 3,900?

Tarun Dua: Yes, yes. So this is CPU, GPU, storage, all capacities put together. So the utilization is closer to 80%.

Deepak Poddar: Closer to 80%. And I think we have around 2,050 lined up in next what 6 to 12 months, right?

Tarun Dua: 2,048 plus some spares that we are planning to deploy in this financial year, starting from May where the first lot of 1024 is expected to go live.

Deepak Poddar: So by FY27 end, you would be targeting around 6,000 of capacity?

Tarun Dua: I would not want to place a limiting number over there, but you could say that that is the minimum number.

Deepak Poddar: At least that. Okay understood.

Tarun Dua: Yes that is what is already visible, so that is what we already spoken about.

Deepak Poddar: Understood. And you mentioned around 80%, 85% utilization by March. I mean, you are talking about by March '27 of this expanded capacity?

Tarun Dua: No. This is the previous gone by March.

Deepak Poddar: Previous March only. Okay, understand. And my second question is on your fixed asset. I mean, we have got around INR1,500 crores order of fixed asset and last two years capex has been around close to that only. So this entire is a depreciable asset? I mean, would that be a fair assumption?

Tarun Dua: I would let Nitin put a perspective on that.

Nitin Jain: Yes, so the entire fixed block is a depreciable asset which constitute of the new B200, which is currently displayed as CWIP.

Deepak Poddar: Correct. And what is the amortization schedule for this? I mean, is it 3 to 4 years?

Nitin Jain: 6 years.

Deepak Poddar: 6 years. Okay. Understood. And what would be FY27 capex target?

Tarun Dua: Like we said we are already planning to deploy 2,048 B200. So that is already in place as far as our known plans are concerned. That being said, we are not limiting ourselves to those plans, so we continue to deploy capital under various business models judiciously to expand the GPU footprint fairly rapidly.

Deepak Poddar: Correct. So any number you have? I mean, I think FY26 was close to INR600 crores, INR700 crores?

Tarun Dua: A look back is more important than putting a number today.

Deepak Poddar: Okay, understood. That would be it from my side. Wish you all the best. Thank you.

Tarun Dua: Thank you.

Moderator: Thank you. Our next question comes from the line of Varun Gandhi from Finavenue Growth Fund. Please go ahead.

Varun Gandhi: Hi. My question is on GPU rental prices and are we seeing pressure on them? I ask this from two aspects. The first aspect is the technology is advancing rapidly. We see the Rubin architecture has been disclosed by NVIDIA which is much more cost efficient for inference compared to your even your Blackwell architecture. And number two is the competition that has been ramping up across domestic competitors. So are we witnessing any sort of pressure on GPU rentals right now or do you foresee that in the near term?

Tarun Dua: See, I think broadly the trend today is that there are not enough GPUs in the world that people want to buy and deploy. So that is the current trend. Now, whether this trend remains for how long is hard to predict, but the broad secular trend has been there that basically there are AI believers and there were AI non-believers.

I think the proportion of non-believers is slowly dwindling and it is quite clear that basically we are barely scratching the surface in terms of utilization of AI. We have looked at so many of our customers, so many of the organizations we work with, including our very own. So where we say, okay, is it the case that our AI utilization is to an extent that we want it today?

The answer for ourselves that we increasingly get is no, this is going to be more like 10x, 20x, 50x of where we are today in terms of how we are utilizing AI. So AI is not coming out of the IT budget. AI is essentially coming out of your budget for the capabilities that you are building for the business.

So that being said, I do not think the demand environment in the foreseeable future is going to be changing negatively for long periods of time. I do not think that is going to be the case from the visibility we have today. We do not see that there is any negative pressure on the pricing today. So, in fact, there are I think good set of tailwinds which are slowly inching up the prices rather than a decrease in pricing.

Varun Gandhi: Got you. So from a demand standpoint, we aren't seeing any sort of negative, but from a technology advancement standpoint, let's say if the Rubin architecture comes in live in this year, and it should, would that pressure?

Tarun Dua: Like there is a lot of workloads that will go into production where they will continue to operate on the targeted architecture. Typically it is what you call funnel kind of a growth where the next layer of the funnel is bigger than the previous layer of the funnel, which means that the demand for existing GPUs will continue to remain very, very strong. So demand for even the Ampere series has remained strong, Hopper series has remained strong.

Varun Gandhi: Tarun, I understand the demand has remained strong, but let's say two years ago, the GPU rental prices for Hopper H200s, are they the same today or they are lower 30%, 40% lower than what they were quoting at 2 years ago?

Tarun Dua: See, I think taking the case of 2 years ago is slightly anomalous from the point of view that I think that was one particular year in which the GPU demand went from say 1 to 10. I think that

capability driven shortages where the capability to print that many GPUs had to be built very rapidly. That created a local maxima of pricing, but broadly we are seeing that there is stability in the pricing for GPUs over an extended periods of time.

Typically we feel that it is possible to definitely use and utilize the GPUs over a 7, 8 year period. We do not think the GPUs are a 2 year story or a 3 year story and then kind of there is something available at a much lower cost and you are able to shift and those shifting costs are very reasonable, so that is not the case. Your entire test debug cycle for doing AI on a new chipset is a whole other set of expenses that you do not want to do that if your job is getting done by a GPU which is available at a very reasonable price.

Varun Gandhi: Got you. Now, my second question is at the beginning of the call, you touched upon how you have significantly improved your software architecture. Could you just give me a brief -- If you could just highlight a brief example of how you've done that? You iterated that you are now trying to capture high-value tokens. Is there some sort of example that you could showcase?

Tarun Dua: So this is broadly from the point of view is that like the tokens themselves today are more valuable. It's just a point of like choosing the right set of focus that people are going to utilize. So, it's not like a -- kind of like a strategy statement to go and capture one particular set of high value tokens today. It's a broad futuristic outlook that like, look, ultimately, we have to go and figure out like which are the high-value tokens.

Now that being said, like from a platform improvement perspective, like we have been working on things like which are required by people running training on larger clusters, people who are running inference how to support them, people who want to build rapidly the pipelines reliably across larger number of GPUs and understanding of like how what is it that the AI data scientists do from a been there done that perspective. So we have made improvements in all those areas for productization of our TIR stack.

Varun Gandhi: Right. So just as an extension to this question, see, we are very well-positioned to capture the sovereign AI tailwind. And my understanding is that BFSI companies would be a key client base over there. Is there any sort of strategic development that you could share where you are partnering with some sort of SaaS companies or other software companies and trying to offer a bundle of software plus compute? Is there something you could help us with?

Tarun Dua: Yes, so as we do something, we will definitely share that with everyone over here. As things happen, like we will do a look-back and say that, okay, this is what we have done.

Varun Gandhi: Okay, but there is nothing that you can share in terms of your thought process over here?

Tarun Dua: See, like we are -- we continue to drown in opportunity. So what we capitalize on, we will come and inform the market.

Varun Gandhi: All right. Thank you very much.

Tarun Dua: Yes, thank you.

- Moderator:** Thank you. Our next question comes from the line of Abhishek from InCred Equities. Please go ahead.
- Tarun Dua:** Yes. Hi, Abhishek.
- Abhishek:** Hi, sir. Hi, sir. Thanks for the opportunity and congrats on a great set of numbers. Sir, my first question is about the guidance commentary. If we summarize your comments on the demand, it appears to be really strong. We have a visibility of 2,048 GPUs and then we said that we do not want to get fixated on the guidance. So just trying to understand that the purpose here was not to get fixated on the number rather than any worries in your mind about the current macro. Is the understanding right?
- Tarun Dua:** Yes, yes. We are not worried about like whatever numbers we give whether we will be able to meet them or not. I think it is like we do not want to give an underwhelming number without first exploring like over week after week over next 52 weeks what we are capable of doing. So it will become sort of like a limiter for ourselves to say that, okay, we are only to meet this number. That is why we are not putting a number today.
- Abhishek:** Perfect. That is very helpful, sir. The second question is on the pricing. So contrary to the prior question, the recent articles are suggesting that the GPU spot rates for April have gone up substantially higher. For an H100, the prices are up anywhere by 25% to 30%. Can you just...
- Tarun Dua:** We are seeing that. We are looking at that. That being said, we -- obviously like we try to capture some of that, but then that being said, we are also trying to support the Indian companies and India as a country, so where we want to have a balance between long-term and short-term. So Short-term of course we want to make more money and long-term of course we want to work with the kind of customers who will sustain and grow on our infrastructure over a long period of time.
- Abhishek:** Perfect. Sir, just last one data point from Nitin sir. So the depreciation for the next quarter should be up by around INR25 crores a quarter, right? Is that assumption right?
- Tarun Dua:** I will let Nitin answer the question around the depreciation. Nitin, over to you.
- Nitin Jain:** Sorry, could you repeat again?
- Abhishek:** Sir, the depreciation for the next quarter should be up by -- it should go up incrementally by INR25 crores a quarter. Is the assumption right, sir?
- Nitin Jain:** Why do you expect that, because?
- Abhishek:** Because of the B200 cluster.
- Nitin Jain:** B200 cluster would be that we are saying the timeline of mid-May. So there would be an increase in that depreciation, but not to that significant effect that what you are telling about.
- Abhishek:** Okay. Perfect, sir. That is very helpful. Thank you for taking my question and best wishes for the next year.

Tarun Dua: Thank you, Abhishek.

Moderator: Thank you. Our next question comes from the line of Neel Munot from Pico Capital. Please go ahead.

Neel Munot: Hi, sir.

Tarun Dua: Hi Neil. Yes.

Neel Munot: Hi. Sir, thank you for the opportunity. So my first question is building on previous participant's question

Moderator: Sorry to interrupt you, Neel, but your voice is breaking.

Neel Munot: Is it better now?

Moderator: Yes, please proceed.

Neel Munot: Yes. So the first question is basically building on previous participant's question. Is that the 80% utilization that we're talking about, I'm assuming this is excluding the new B200s that have come in?

Tarun Dua: Yes, so that is excluding the B200, obviously.

Neel Munot: Yes. So when we talk about the 80% utilization, with respect to the remaining capacity available, have we kept these aside for POCs or demand is still ramping up in that sense?

Tarun Dua: No, so the demand is still ramping up over there. It is like all sorts of different sets of capacities. So, like, they all work in different combinations. So an ideal utilization number could be, it depends on basically where do you end up. But that is not a limiting factor on the MRR from a very straight line perspective. So for instance, like the same particular piece of hardware could under different circumstances be billed at wildly varying rates.

Neel Munot: Sure, my assumption...

Tarun Dua: These are all averages essentially, so these are all important from a look-back perspective. So this is our assumption on that, look, we are like doing what is sort of an 80% capability. Now 80 could be 80, 80 could be maybe 85, 80 could be even 75, maybe even 70.

Neel Munot: Okay, understood. And sir, with respect to the demand scenario in B200s, have we closed any -
- AI Mission?

Tarun Dua: We have not closed any deals on the Blackwell capacity which is going up as yet. Obviously once we close a deal, we will kind of inform everyone over here.

Neel Munot: So, the reason for this is, is the reason ramp-up or is the reason that we're still closing POCs or in that sense?

- Tarun Dua:** There is a lot of interest in the Blackwell capacity that we are building. We are working towards with multiple customers to figure out what sort of problems they need us to solve for them and what it is that they are looking at from us, and we obviously want to work with the customers that are best placed to utilize what we have to offer.
- Neil Munot:** And sir, one more...
- Moderator:** Neil, please rejoin the queue if you have more questions. Thank you.
- Neil Munot:** Okay, sure.
- Moderator:** Thank you. Our next question comes from the line of Hardik Gandhi from HPMG Share and Securities. Please go ahead.
- Hardik Gandhi:** Hello, hi sir. Thank you for the opportunity. Congratulations on a set of good numbers.
- Tarun Dua:** Thank you, Hardik.
- Hardik Gandhi:** I just wanted to push on a question asked by the previous participant. I think so majority of the investor group just wants to know whether the revenue uptick is due to any lumpy nature or are you seeing any sustainable growth in the long run. Because in the last December and March we had, we suddenly dropped the numbers. We stated that the startups were no longer wanting our ecosystem, right? So we just want to know...
- Tarun Dua:** Sure, sure, sure. Let me try to answer that question as best as possible. See, the lumpiness is a part of the nature of the business. The lumpiness will keep going down as we keep increasing our GPU base. That being said, we are currently seeing the kind of customers who are buying these GPUs have sustainable long-term workloads.
- These are not of the short-term nature for the most part. We do expect to have sustainability for these workloads over the medium term. I do not think there is going to be a massive amount of lumpiness in the near term. Medium term of course, as the number of GPUs grows, I think the effect of lumpiness would be far more muted than it has shown up in the past in the early days of us building the capacity.
- Hardik Gandhi:** Right, right. And the second question from my side was to understand again on this part itself that what percentage of revenue are we allocating to long-term contracts so that our MRR doesn't...
- Tarun Dua:** It is of a dynamic nature. Basically, it is not possible to predict...
- Hardik Gandhi:** Yes, yes. We are not predicting. I am just saying from a cost...
- Tarun Dua:** We are not setting any hard lines over there. See, we are not setting any hard lines over there. There are I think far more number of variables to consider than it would be possible to describe over here that what to pick up when in terms of customer interest.

I do not think I would be able to give a very fair fixed answer to that question. We are definitely interested in having some percentage of long-term contracts. What that percentage will come out to be is something we need to look at on a look-back basis when the size and scale of the GPU installation is sufficient. I think in the short term, there is no clear-cut answer to that.

Hardik Gandhi: Right. Sir, just a small data point if I can. What percentage of revenue did come from customers outside India versus domestic demand for the last quarter?

Tarun Dua: I will let Nitin take a stab at this question.

Nitin Jain: So, the international revenue for the last quarter is roughly around 35 odd percent. 35% to 37% is the international customer revenue.

Hardik Gandhi: Okay. Yes. Thank you, thank you. That is really helpful. Appreciate the time.

Moderator: Thank you. Our next question comes from the line of Srinivasu K from Trust Investment Advisors Private Limited. Please go ahead.

Tarun Dua: Hi Srinivasu.

Srinivasu K: Hi, sir. Am I audible?

Tarun Dua: Yes, yes. Loud and clear.

Srinivasu K: Yes. So, sir, what percentage of Q4 revenue came from inference workloads versus training and any future split up of this mix at least directionally, sir?

Tarun Dua: So, I think there is a lot of confusion about what consists of inference revenue versus training revenue. I think these are not super hard lines. Sometimes the same set of GPUs are used by the customer for kind of bursting up the inference workloads when the inference workloads are high, and at a time when the inference workloads are low, they get converted to training.

So, the workload management is far more flexible today to point into what percentage of inference, what percentage of training today. And I think again, at a much larger scale, these numbers would become clearer.

Today of course, it is hardly possible to super differentiate between who is using that one GPU, four GPU, eight GPU, sixteen GPU cluster, either for training or for inference. But that being said, majority of revenue is still closer to training than to inference.

Srinivasu K: Okay, sir. And you said 6,000 GPUs by FY27, at least minimum, right? So, what would be the upside, sir? Could you deploy 8,000 on upside, 8,000 to 10,000?

Tarun Dua: I do not want to put any numbers right now. Let us look at those numbers as a look-back over the coming period of time rather than putting out a number today.

Srinivasu K: Okay. Thank you and best of luck.

Tarun Dua: Thank you, Srinivas.

Moderator: Thank you. Our next question comes from the line of Rohan Nagpal from Helios Capital Management. Please go ahead.

Tarun Dua: Hi Rohan.

Rohan Nagpal: Hi. Thanks for the opportunity and congrats on the good set of results. My first question is, you said that there is a variety in terms of your current MRR being short-term, medium-term, slightly long-term. Could you provide some directional commentary on how much of that MRR is 6 months plus or 1 year plus?

Tarun Dua: Yes, so I think one is we are still on a very small base of GPUs. We are not doing any hard split in terms of where we want to be or where we are today. It is fairly dynamic and rapidly changing, both with new capacity coming up online. So all these numbers are subject to very rapid and very wide amount of change. So these do not add to any understanding for anyone today.

Rohan Nagpal: I mean, if there is a 20% to 25% increase in the spot price, then capacity that is contracted out for 6 months or 1 year will not be in a position to take advantage of any increase in realization, right? So I think it would be somewhat beneficial for someone looking at it from the outside to get a sense of some amount of revenue uplift or exposure to increased realizations.

Tarun Dua: So that is a very granular set of understanding that you are seeking on a very small size of the infrastructure installation base compared to the global peers. I think we would like to avoid this conversation for today. As we grow in size, then we kind of talk more about that, okay, now these numbers actually mean something and they are making sense to someone on the street. Today these numbers would not make much sense.

Rohan Nagpal: Okay, sure. Thanks for that. And the second question was, so you talked about higher value tokens. Could you provide some color on which industries are sort of taking advantage of higher value tokens from your vantage point?

Tarun Dua: Not really. I think it is more closer home than that in terms of what is a higher value token. Asking a question like, okay, what is closer to a search query is a reasonably low-value token. Someone using the tokens to solve a specific problem that is paying the bills for a company is a completely different value conception for those tokens.

So answering a basic transactional support query is a reasonably low-value token. Being able to answer a more complicated query or being able to confirm a more complicated transaction with a customer on call is a higher value token. There are infinite number of variations of how people use AI.

The key is to find those customers who are utilizing the higher value tokens where they are not worried about, okay, what is my ultimate per hour GPU price, but what is the cost of my solving a particular problem and generating a higher ROI. Again, we are not directly involved all the time in identification of higher value tokens and kind of being able to strategize in a very, very

static manner that oh, we should go and sell to this customer, they will probably pay higher for this particular GPU.

It is a broad industry trend where we are seeing increasingly the demand coming from customers who are not price sensitive because of their ability to generate or consume higher value tokens through a combination of their own software or proprietary or open source software that is available extant in the market today.

Rohan Nagpal: So in that case, what is E2E's role in this higher value token generation? I mean that would just be an outcome of if the person who is renting a GPU or taking on GPU capacity is able to extract higher value they will be willing to pay more, right? So the price discovery should just take care of who is generating higher value token. I was not too clear on how E2E is sort of working towards higher value?

Tarun Dua: See, how you are able to quickly get to that point of being able to generate and utilize those tokens for your business. It is fundamentally how to build on various parts of the software and hardware stack to make it available quicker. I think that is where the role is played by E2E.

Rohan Nagpal: Understood. Okay, fair enough. That is it from my side. Thank you very much.

Tarun Dua: Thank you.

Moderator: Thank you. Ladies and gentlemen, due to time constraint, that was the last question for today. I would like to hand the conference over to the management for the closing remarks. Over to you team.

Tarun Dua: Yes. Thank you to all our investors, all our team members, all our customers, all our ecosystem partners for all the support you have all extended to E2E Networks over last many years. It has been an overwhelming amount of love for us that we are seeing in the community and overall community, and we wish you all a great day ahead, and thank you for patiently listening to our call. We would like to kind of once more thank everyone and end this call.

Moderator: Thank you, sir. Ladies and gentlemen, on behalf of E2E Networks Limited, that concludes this conference. Thank you for joining us and you may now disconnect your lines.