



## **E2E Networks Limited**

CIN- L72900DL2009PLC341980

1st Floor, A-24/9, Mohan Cooperative Industrial Estate

Mathura Road, New Delhi-110044, Phone No. +91-11-4084-4964

**Email :** [cs@e2enetworks.com](mailto:cs@e2enetworks.com), **website** <https://www.e2enetworks.com/>

**Date:** November 4, 2024

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 Q2/H1-FY25**

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, October 28, 2024 at 11:30 A.M. (IST) to discuss the Company's Q2/H1-FY25 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 Gaba**  
**Company secretary & Compliance Officer**  
**ICSI M. No.: A59215**



**“E2E Networks Limited**  
**Q2 FY '25 Earnings Conference Call”**  
**October 28, 2024**

**MANAGEMENT: MR. TARUN DUA – MANAGING DIRECTOR – E2E  
NETWORKS LIMITED**  
**MS. MEGHA RAHEJA – CHIEF FINANCIAL OFFICER –  
E2E NETWORKS LIMITED**  
**MR. RONIT GABA – COMPANY SECRETARY – E2E  
NETWORKS LIMITED**

**MODERATOR: MS. SOUMYA CHHAJED – GO INDIA ADVISORS**

**Moderator:** Ladies and gentlemen, good day, and welcome to E2E Networks Limited Q2 FY '25 Earnings Conference Call, hosted by Go India Advisors. 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 the star, then zero on your touchtone phone. Please note that this conference is being recorded.

I now hand the conference over to Ms. Soumya Chhajed from Go India Advisors. Thank you, and over to you, ma'am.

**Soumya Chhajed:** Thank you, Steve. Good morning, everyone. We welcome you to E2E Networks Limited Q2 FY '25 results con-call. We have with us on call today Mr. Tarun Dua, the Managing Director; Ms. Megha Raheja, the Chief Financial Officer; and Mr. Ronit Gaba, the Company Secretary. I must remind you that the discussion on 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 now request you, Mr. Tarun Dua sir, to take us through the company's business and highlights, subsequent to which we will be opening the floor for Q&A. Thank you, and over to you, sir.

**Tarun Dua:** Thank you, Soumya, and good morning everyone who has joined our call today. I would like to welcome all of you to E2E Networks Q2 earnings call for the financial year ending March '25.

Let me briefly reintroduce you to the company. E2E Networks is a cloud computing company majorly focusing on accelerated computing using cloud GPUs and was founded in 2009. Today, we support a wide range of customers ranging from startups, enterprises, higher education and research. And we are also MeitY empanelled, which means that we are also targeting government for our solutions.

We have been an early mover into the space of AI, ML, and GenAI. We have operated our cloud GPU platform. The very first GPUs were bought in 2018, but we have operated our cloud GPU platform from all the way from like 2019 on our cloud. And E2E continues to leverage the early mover advantage by focusing on the three areas of support, solutioning and software. So we have been supporting our customers, including many of the unicorn scale customers since 2009.

We have provided AI, ML solutioning for our customers since 2019 onward. And we continue to deliver excellent software, which matches up to some of the global majors in terms of AI, ML, and GenAI computing via our TIR platform, and this platform was released like in the last year or so.

So our core services and platforms include ability to support an entire cloud platform, which typically provides all the major features that are required for majority of the bill of materials of any cloud user in India. And our GPU platform, like we try to compare it with the best of best in terms of being able to do rapid deployments of like training clusters, inference workloads, and model endpoint deployment and any kind of like GenAI support. We plan to include that in the coming future.

So we have a deep and vast focus today on AI, ML, and GPU computing. Now, as a world kind of like additionally add GPU computing apart from the traditional CPU-based cloud systems, we believe that like we are very well positioned to leverage many of those workloads on our cloud platform and meet the demand for GPU computing. We have explained that in the previous call as well, so a typical CPU today would be say about 250 to 300 cores, a typical CPU server would have about 250 to 300 physical cores.

A typical GPU server, let's say an H100 or H200 would have 17,000, nearly 17,000 into eight physical CUDA cores, which essentially means at about 10x the cost and 10x the opex compared to a typical CPU server. Now, how it matters to us is that like this obviously, on a per teraflop basis is like cheaper compute.

And this opens up the world of possibilities for the typical CPU using organization to kind of like who are only processing say a couple of percentage points of the organization's structured data to start dealing with like training, like a lot of data in motion, the unstructured data through the AI/ML workloads that run on top of GPU, they can start using that.

And for us, like obviously, the ARPU grows, because essentially we are providing like typically to a typical customer like amount 10x or 20x the amount of ARPU for about 100x to 200x the amount of compute, essentially because like compute on GPU, cloud GPU is cheaper than say cloud CPUs, because CPUs tend to have like a lot more cores.

Now, GPUs, of course, need specialized software to operate. And in the recent past, we've seen that the open source state of the art has rapidly evolved for the customers to be able to leverage all of this compute, which is now available to be able to process the unstructured data.

Now, the source of competitiveness in the AI/ML and GPU space comes not just from the sticker price of cloud GPUs that we offer. But the main source of cost efficiency and competitive advantage comes from the fact that ultimately, like if you have a large number of GPUs where you're spending significant amounts of money, like if you are able to prevent the idleness of those cloud GPUs by building up the platform software in a way that people are able to rapidly deploy clusters, as opposed to spending a couple of days to build out the workload managers onto a cluster. So that results in significant deployment savings.

Now, whenever you are putting up new workloads, again, like all of that setup needs to be done again. So if your setup allows you to easily tear down and rebuild the software stacks required for cluster-based training or deployment of scale to zero inference workload, essentially, you save a lot of time and effort from the engineering side.

And apart from saving the engineering time, you also end up optimizing the amount of time your GPU stay idle versus -- you are able to reduce the idle time and you are able to utilize like most of the GPU time that you have purchased.

Typically, with H100 onwards, we have seen that there is less usage of hourly GPU time compared to the past in the A100 era. So H100 and H200 onwards, like what we've seen is that like the contract length is typically one month to three years, which means that like you have already reserved the capacity, which means that the software that you have in terms of being

able to schedule, being able to set up the clusters, you are rapidly able to do what you want to do, which means that the movement of data, the creation of operating system environment, the setup of the configuration, all of that has to be supported by the software in a way that like all of these activities get done faster than your typical raw GPU providing competition.

So that's where the value of the platform really comes to the fore. So competitive advantage just doesn't come from like being able to offer at a lower per hour sticker price, but also the efficiency of the efficiency and the usable-ness of the software that you build on the platform.

Now, we have reiterated in the past that the growth opportunity for India is very, very clear. India generates like more than 20% of world's data, thanks to our vast and digitized population. And currently, in the data centers, like we, if you just measure by power, we are probably like somewhere around 2% of world's compute capacity, give or take something. So this highlights a very major opportunity in terms of expanding the compute capabilities, especially for accelerated computing led by cloud GPUs in our part of the world.

Now, essentially, if you look at the past, there were companies who were doing their own computing. And there were companies who were using like large hyperscalers to do the computing, so large multinational hyperscalers. And the middle ground of people who had built like a cheaper cloud, cheaper or more value-centric cloud at a lower price point, like that was a kind of like a much smaller space.

Now, more recently in the cloud GPU, as the cloud GPU market has evolved, in the other developed markets, we have already seen that there is a middle ground between having to run everything on your own. And that is especially hard with regards to like running very large GPU focused machines, which are very expensive, and the amount of time it takes to set up is all adds to your cost.

So the middle ground is that like you take up providers in India, like cloud GPU providers in India, like E2E, who provide a middle ground in terms of cost. And in terms of capabilities required to operate, compared to either running your own where the cost is low, but the capabilities required is very, very high.

Or going to hyperscalers where the capabilities required to run are reasonably low, but the cost is very, very high. So E2E offers a middle ground where like a reasonable set of capabilities allows you to use our platform very, very easily, and recover most of the cost advantage that you had by doing it on your own.

Now, this is also the cloud GPU has also accelerated the trend of cloud repatriation, where people have run through multiple cycles of running their things, running their workloads on the cloud. And they have found that like, yes, the initial calculations on the cloud, on the hyperscaler clouds typically tend not to hold, and the cost rapidly rises. And this is where players like E2E who offer a very, very predictable pricing with not too many variables doesn't require an advanced degree in mathematics to operate and understand that where the pricing could move in the future.

So we come into the picture over here. Now, what we have seen, like more recently is that like, the demand for cloud GPU services continues to be fairly robust. And the cloud GPU utilization

rates like although they have come down from the last quarter or so with the new addition of the capacity very recently in October, but we continue to see a reasonable demand pipeline. And we hope to achieve higher and higher rates of GPU utilization. And we are seeing a lot of interest in terms of the adoption.

So I think in the past couple of years, we have seen like startups, of course, are the first ones to adopt. AI startups are obviously leading the pack in terms of adoption. But increasingly, we are seeing like budget allocations from enterprises with regards to figuring out like what are the core, some of the core problems, identification of some of the core problems that they see, which can be solved using AI, generative AI and machine learning.

So there is a very strong impetus towards like, initially, trying out, figuring out, putting through the blue board to figure out the AI solution, and demonstrating that to the decision makers within the enterprises.

So there is a very, very strong push towards that. And like whatever, kind of like in the enterprises, like a lot of small language models running like the previous generation of GPUs are actually seeing like production deployments today. So this is something that started two years back. So which means that we are going to see the push towards like using more cloud GPUs in production spread over the next many years. So where this trend has already started.

Now, kind of like latching on to this trend of like some very strong demand and growth coming in this field. We, E2E has kind of like evolved its infrastructure and expanded it. We have nearly something like 700 or so 800 GPUs. We have recently secured 256 H100 GPUs in October. And we also continue to expand the non-H100 or non-Hopper variety of GPUs that includes the T-force and the L-force and L-40s and A-100s and A-40s of course, like we are only expanding the latest generation of GPUs.

So currently that number also as of the end of September stands at about like 600 or something. And your company has also raised nearly INR405 crores, which would be used primarily for infrastructure expansion and further investing into in-house AIML research and GenAI research, which would result in outcomes that can be utilized in our platform to help our customers do things faster, cheaper and better.

So like the market potential as per various industry reports and industry bodies seems to be presenting massive growth opportunities in this area of both the cloud and upward from the stack. And in India alone, like the projections are running into the amount of value AI would provide to the GDP would exceed like a couple of hundred billion dollars. The AI infrastructure market growth is projected at a reasonable CAGR of like 25%-30% year-after-year, which increases the demand for cloud GPUs for selling the AI applications.

So, E2E aims to support the AI-driven business transformation in all sorts of organizations, whether it is startups in India, startups outside India, higher education and research. We want to work with the government. We want to work with enterprises. We want to work with multinational organizations who have their centers of excellence in India.

So, we continue to build solutions to support various industry solutions to support all -- a diverse range of customer base. And with this note, I would like to hand over the call to Ms. Megha, our CFO, to give a brief update on what has been happening in the last quarter and overall growth that we have experienced. So, over to you, Megha.

**Megha Raheja:**

Thank you, Tarun, and good morning, everyone. I would like to start by highlighting some key financial metrics for our performance in Q2 FY '25. For Q2 FY '25, total revenue was INR484 million, reflecting a substantial year-on-year growth of around 120%. Our EBITDA for the quarter stood at INR314 million, showcasing an impressive growth of 181% year-on-year, with an EBITDA margin of 66.1%, an increase of 1,440 basis points compared to last year. In terms of net profit, we reported PAT of INR121 million, demonstrating a remarkable year-on-year growth of 108%.

The PAT margin for September quarter was 25%, and our diluted EPS for quarter is INR7.8, marking around 98% increase year-on-year. In comparison to previous quarter, June FY25, we experienced a total revenue growth of 16%. Our EBITDA also grew by 15% quarter-on-quarter, resulting in net profit of INR121 million, which reflects a quarter-on-quarter increase of 18.9%.

Additionally, I am pleased to announce that we successfully raised INR4,056 million through preferential issue of equity shares, which position us well for future growth and investment opportunities. Pursuant to this, the paid-up capital has increased from INR145 million to INR169 million, and our total equity has increased from INR718 million to INR4,874 million as on September 30, 2024. That concludes financial update for the quarter.

**Tarun Dua:**

Megha, I think you mentioned INR405 million, that is INR405 crores.

**Megha Raheja:**

INR405 crores, INR4,056 million.

**Tarun Dua:**

Okay, sure.

**Megha Raheja:**

Okay, with this we can open the floor for question and answer.

**Moderator:**

Thank you very much. We will now begin the question-and-answer session. The first question is from the line of Ritesh Chadha from Lucky Investments. Please go ahead.

**Ritesh Chadha:**

Thank you for the opportunity, sir. Sir, one broader question I have. I understood the growth possibilities in India and your initial comment on the GPU capabilities. With this INR400 crores of capital raised now that we have on our balance sheet, and last year the computer's gross block was about INR222 crores and the total block of asset was INR260 crores. From a three-year perspective, let's say FY '25, FY '26 and FY '27, what should be the asset creation or the asset capex that we will do over the next three years and this incremental capex that we are doing and what asset turn and margin this capital will be invested?

**Tarun Dua:**

Hi, Ritesh ji. Thank you for your question, Ritesh ji. So essentially, as we have said many times in the past in our conference call and earnings call, mostly we recommend that you look back into the past for many of these answers. And for the future, like it's very, very hard to predict, like it's a growing market.

So we have no upper limits in our mind with regard to how the growth block should grow. It completely is driven by the demand. So we are able to like set up our deployment in a way that like we are prepared for large capacity expansion.

But obviously, we keep building the capacity as the demand keeps coming in and as the customers keep getting signed up onto our platform. So based on that velocity, we keep on building the infrastructure. Now, obviously, there is some sense to this fundraise.

So we are essentially in a way we feel that like this should be something close by in terms of like how quickly we can ramp up and deploy. So it gives a sense of perspective. The perspective for asset turn and EBITDA margins, of course, like we have mentioned that many times that like the EBITDA margins are also a function of the platform effects that we are seeing.

So obviously, we have certain fixed costs with regard to having an employee base and certain amount of investments we make for the data center side like they tend to spread across like in the sense that yes, for the most part, it is pay as you go, but then you also resolve certain fixed amount of capacity or promise certain fixed amount of capacity to be paid for.

So which means that all in all, like as the volume grows, the EBITDA margins like typically tend to grow. Now, we are obviously hopeful of further increasing the EBITDA margin in the coming years.

Now with regard to kind of like how the growth stands out, like again, like we are very, very hopeful based on all the industry reports and the expansion of industry and what we are seeing in terms of the interest level coming in from like a lot of people to figure out like how to do how to use the Gen AI to advance further their organization's goal.

So we do hope to grow better than the industry average being a much smaller company. So that is what I would like to point to.

**Ritesh Chadha:** Sir, at least can you share one, you might have a goalpost. So at least on the gross block side, if you could share that over the next three years, I would like to have, let us say, INR1,500 crores worth of assets, which will be given out?

**Tarun Dua:** Obviously, it is a function of demand. So if you look at like the amount of fundraise and the ability to leverage that, so it could be anywhere between say INR300 crores-INR400 crores of new addition in terms of gross block all the way to like maybe something like INR1,000 crores to INR1,500 crores. So there is no upper limit in our mind that like it cannot go beyond like any number. And obviously, it is completely based on like the demand we are able to capture.

**Ritesh Chadha:** Okay, can I ask you in a different way for FY '25, the current year for which you will know the visibility, what will be the assets that you will add for the current year?

**Tarun Dua:** Okay. So again, I feel that like it is better to look at like the kind of like... So assets like it completely depends on like a number of factors. So one is, of course, like there is a slight thing about like supply chain. So it takes like a couple of months for the supply chain to provide us



the new material. Second, of course, that like we always look at like, okay, what's the current pipeline and what's the current capacity.

So based on these numbers, like we take all these decisions. Now, typically, what we are doing today is like we are buying hardware in blocks of about approximately INR100 crores plus minus INR10 crores. So...

**Ritesh Chadha:** INR100 crores per annum or INR100 crores per quarter?

**Tarun Dua:** Every time we buy, it's like a typically like INR190 crores to INR110 crores or INR120 crores kind of a buy for the hardware. And we keep on trying to run multiple cycles of this by like every year. So, basically based on that like it's very hard to predict. You could run more number of cycles, you could run less number of cycles. It depends on the demand as well as the supply. So, we are quite flexible in that regard. So, I wouldn't like to put a number to like what happens this financial year, what happens next financial year. Like, we always recommend that in our cloud business like look back rather than look ahead.

**Ritesh Chadha:** Look back. So, if I look back then you added INR175 crores last year.

**Megha Raheja:** No, we added INR185 crores last year in FY '24. And in H1, we have added around INR107 crores till September.

**Ritesh Chadha:** And my last question is at what is the comfortable balance sheet leverage for you and the GPU-CPU mix for us in the asset, what is the GPU-CPU mix today and what will be the GPU-CPU mix in the future?

**Tarun Dua:** The GPUs in the mix keep growing in number in terms of percentage. So, I think like we are already at like more or less like 90-10 kind of a mix. So, I think that's the steady state in terms of what we buy the GPU and CPU mix. So, with regards to I think like comfort in terms of leverage, I think like that depends on the MRR growth for us.

So, with more MRR growth, we become more comfortable of being able to service more debt. So, again like as far as like debt is concerned, leverage is concerned like there is no upper limit in our mind. As long as we have the ability to surface it comfortably, we can take more debt.

**Ritesh Chadha:** Okay. And lastly what is the asset turn that you would look at for deploying your capex, the asset turn ratio? Is it half time, one time?

**Tarun Dua:** Currently, it's like broadly about 0.5 for incremental deployment. So, like we of course like as we continue to build our software like we expect it to grow. Now, what phase it grows is very hard to predict. So, which is why we say that like, yes, like generally like let's look at what is currently happening. And when it changes like, for example, like when it changes then we know that it has changed.

**Ritesh Chadha:** Okay. And last sir you do not have any working capital...

**Moderator:** Sir, I am sorry to interrupt Mr. Ritesh. Could you please go back in the question queue.

**Analyst:** Okay.

**Moderator:** Thank you. The next question is from the line of Ketan from Taurus Investments. Please go ahead.

**Ketan:** Hello, sir. Can you please elaborate the MOU signed with Dell the prospects whatever it is?

**Tarun Dua:** So, see Dell is one of the industry leaders in terms of providing a lot of hardware solutions in the entire world. And that gives them a unique access through their entire sales and marketing organization into what is happening across the organizations across the globe. And what are the trends that we are seeing in India which could accelerate based on their assessment of what is happening in the rest of the world.

Now, at one level, of course, they are a vendor to us where they are one of the vendors who supplies the GPU servers and other hardware to us. At another level like essentially both our companies are in the business of providing solutions to our customers. So, in that essence like from a sales perspective if we are able to jointly target solutioning for a few customers, then that works out well for both us as well as Dell.

So, that's the broad nature of the tie-up that like ultimately we are looking to kind of like help each other serve our customers better. So, since Dell does not provide the cloud, so they are able to kind of like point to us towards like some of the cloud solutions for the customers that are asking for solutions from them. And in our case like in case like the customer is not amenable to a cloud, then obviously kind of like we can give them the options to directly work with Dell.

Obviously, we are not in the business of selling hardware. So, that's broadly it. So, it's like very early days. We have not seen any revenue impact because of this MOU or this relationship as of today, but obviously we kind of like continue to work towards helping each other jointly in the cloud GPU and the GPU market in India.

**Ketan:** And what is the latest MRR?

**Tarun Dua:** Sorry, say that again please.

**Ketan:** Your latest MRR month revenue rate?

**Tarun Dua:** Yes. So, we have updated the presentation in which the exit MRR for September is mentioned So, I think it is somewhere close to 16.5 something.

**Megha Raheja:** Yes 165 million.

**Moderator:** Thank you. The next question is from the line of Pankit Shah from Dinero Wealth. Please go ahead.

**Pankit Shah:** Hi, Tarunji. Thank you for the opportunity. Sorry for harping on the capex front again. So, actually if you remember in the earlier call the plan was to do like a capex of approx INR800 crores in a scenario where things goes as per the plan including the fundraise also. So, now if

you are done with the fundraise, should we expect that things are as per the plan, as in we should do 800 or more?

**Tarun Dua:** See without putting a time limit to it I think like over the next couple of years we should definitely do more capex than INR800 crores.

**Pankit Shah:** That is for sure. But I'm saying from a visibility perspective as in and also in terms of leverage?

**Tarun Dua:** Pankaj ji, I already answered this question. So, where essentially what we are saying is that like every time we go and make a hardware expense typically it goes something like this. Let's say essentially we are buying say somewhere close to 255 GPUs, associated storage, associated memory, associated peripherals, some amount of CPU compute and that entire bundle would typically like in a one series of purchases, which are close to each other to build capacity. So, that would be close to about like somewhere between INR90 crores to INR120 crores or maybe INR125 crores.

And typically, we try to kind of like consume majority of this capacity and see the visibility for the next set of capacity before we go and place another order. So, essentially it's a balance between building the capacity versus the run rate for sale. So, ideally we would like to run as many cycles as quickly as possible. Now, that being said like there would obviously be periods like for example, like the holiday season due to festivals where things would get delayed by a couple of weeks.

But broadly, the idea is to run as many cycles of these throughout the year to build up the capacity. So, again very hard to predict like what numbers do we end up with, but essentially we are moving as fast as possible.

**Pankit Shah:** But if I want to understand, say, currently, in current times, how much time does this one cycle consume, say from ordering deployment to say consumption, if you can give a rough ballpark figure of two months, three months or?

**Tarun Dua:** Okay. So, this keeps on changing, like the hardware supply chains, like sometimes they become very efficient. Sometimes there is a bit of inefficiency in those hardware supply chains. Then the conversion of demand, demand is a constant in the sense that like, if you look at the amount of workloads being run currently in India on the cloud, those number of workloads is essentially infinite.

Like, so, we are talking about a couple of billion dollars worth of workload spread across CPU and GPU, where obviously, the workloads on the GPU side are growing very, very rapidly. Now, so one is, of course, like the demand pipeline coming from the presence of that many workloads being run by that many enterprises and organizations.

And the second part is the conversion velocity. So, again, like I said, that like the conversion velocity is dependent on like a number of factors, including solutioning for individual customers, depending on their size and scale. And kind of like, it's very, very hard to predict all of these numbers. So, instead of like predicting, we react. So, that is what I would like to again stress on that, like, we work very nimbly and efficiently.

And kind of like, you see those efficiencies, I hope that we are able to show, reflect that efficiency in our numbers as well. So, we intend to be the same in the future to kind of like, nimbly react to the market conditions very, very rapidly, and not commit to like hardbound plans. Because ultimately, what happens is that like the hardware SPUs, including cloud GPU SPUs change very rapidly.

So, it is very important to build only the amount of capacity that you can kind of like sell very rapidly and quickly, and kind of like deploy more towards the future version. So, this is what reflects in our, this reflection of our philosophy also shows up in our, us being the first ones to deploy H200s in India. So, like, this flexibility goes away, like if we go and commit ourselves to a hardbound number. So, that's what I would like to continuously maintain over here.

**Pankit Shah:**

My second question, wanted to understand that, how is the competition shaping and with a lot of H100 coming in, say, inventory from competition side, we are seeing any impact on the yield side or anything on that end?

**Tarun Dua:**

So, as yet, we have not seen like a major impact of competition in terms of the yield. So, with regard to like the overall scenario of competition, I think it's also a reflection of the growing market. Now, some of the largest companies of the world have already been working in this space.

And we do see that like, both from the end of like the bigger companies as well as from the end of the smaller companies, like, there would be more participants. I think we have mentioned that like, this is such a large market that it would not be kind of like impossible to imagine that there could be like 10 successful players, each with its own way of like operating in this market and its own niche could be successful in India alone.

And a few of them would be of course, multinationals, a few of them would be domestic players. So, that's what it looks like today based on the various industry moves that we are seeing in the cloud GPU market today. So, there is obviously growth of competition, as there is growth of the market. So, I think both go hand-in-hand, if it is a very attractive market, obviously, there would be competitors.

Now, we of course, have our own USPs in terms of having done support for more than a decade for unicorn scale customers, having done solutioning on the cloud GPU side for more than four years plus, and having built our software over the past couple of years, which is being better tested by a lot of our customers in terms of adoption. So, which obviously gives us an early mover advantage in terms of having a better understanding of what the customer needs are.

And with regards to yields, like we also try to limit our exposure to any one FPU more than by going towards the newer FPU as soon as possible. So, especially by being better at predicting that okay, what is the volume that we will be able to sell on an FPU. And later of course, like for next couple of years, that FPU stops getting new supplies because it stops getting manufactured.

So, eventually, the price for any FPU reaches a steady state. And that steady state like eventually is still a reasonable number. This is what we have seen in the past with the previous generations

of FPUs. So, where initially there is less demand for an FPU because people have not yet figured out like how to utilize this FPU with the newer software library. And then over a period, the demand increases, the supply increases, then eventually the supply drives up. And this particular price point for this particular FPU then kind of like reaches a steady state.

And then that steady state continues for next many years. So, this is something we have seen in the past. Now, how it exactly pans out for all future FPUs, again, is yet to be seen. But like looking at the past, this is what we feel is likely to happen.

**Moderator:** Thank you. The next question is from the line of Ashish Sriram from GM Financial Mutual Funds. Please go ahead.

**Ashish Sriram:** So, how do you see your collaboration with People+ai in terms of large scale computing? If you could help us understand that.

**Tarun Dua:** See there are many organizations which do not have like a kind of like very hard profit motive. They are trying to figure out a play for India. So, I think like People+ai is one such organization. Of course, like we are working with many others as well, where they are trying to figure out that, okay, from an India perspective, like, what is it that makes sense for like all of the cloud FPU players and maybe even other CPU cloud players in India to kind of like jointly do together.

So, like, it's still very early days. It's again, like, so there is no current revenue impact due to People+ai in terms of a positive or negative in any way currently. But we are hopeful that, like, as People+ai evolves its strategy for where they want to influence the direction of the country, like, it will become known that, okay, what's the impact they're able to generate?

**Ashish Sriram:** So, in terms of supply chain, do you still feel that we are still restricted in certain way in terms of sourcing of GPUs? And in parallel to that, how do you see the government's AI mission panning out?

**Tarun Dua:** It's a very complicated answer. Both yes and no. So, in the sense that, like, we are definitely, India is definitely not a top buyer of GPUs. And so to speak, like, we do see kind of like India slightly behind in the queue from people who are able to kind of like get the very first versions of the GPUs that come out. Of course, like, even hardware has like a set of bug fixes that can be applied via the firmware and so on.

So initially, like the very first people to get, like, the initial samples of GPUs tend to be the largest of hyperscalers or the largest of consumers of GPUs who are buying 100,000 plus kind of GPUs a year. That being said, that is still not a sufficient quantity for them to be able to launch like a broad based service, on demand service for their entire customer base. Let's admit, we do see that, broadly, the bulk of the GPU supplies, almost everyone gets almost at the same time.

And so the current version, for instance, I think, like, it's now available for anyone to be able to kind of like purchase. And of course, the next version, like we intend to place the orders much in advance to ensure that, like, as soon as the new supply comes in, we are the first ones in the queue. So yes, like the supply chain, once a new SKU has stabilized, it's like very, very fast.

When you're looking at a new SKU, there is like certainly a wait period. Now, a lot of things have improved compared to what they were a year ago, or even 6 months ago, in terms of speed of delivery. And the entire project being like shaped together. So for example, when you're buying GPU servers, you're not just buying the servers, you're also buying the networking equipment, you're also buying the fiber optic cables, you're also buying the storage, you're also buying the memory.

So like, I think all of that, as of today, currently, it's all in a very good shape from a supply chain perspective. Now, there could be, like always, destabilization of a few components which are critical to the deployment, like in the future, as well as they have been in the past. But then, as of today, we don't see any impact on the current SKU supply chain.

**Ashish Sriram:** Yes, fair enough. So, let me talk about...

**Moderator:** Sorry to interrupt, Mr. Ashish. Can you please fall back in the question?

**Ashish Sriram:** Sure.

**Moderator:** Thank you. The next question is from the line of Shankar from Singularity AMC. Please go ahead.

**Shankar:** Yes. Hi, Tarun. Congrats again on the numbers and the performance. Would you be able to talk about the progress made, if any, with the partnership? And also, are you seeing most of the demand coming from domestic startups or international? If you could touch upon...

**Tarun Dua:** The split between domestic and international has been about, I think, like close to 65, 35, approximately, in the last two quarters, is what we have seen. And mainly from outside, we see a lot of interest from the startups and some higher education and research institutions as well. And from a making progress perspective, like, again, it's a long cycle business.

So, we kind of, like, continue to figure out, how well we are able to operate in the government ecosystem. So, it is still very, very early days. So, as something major becomes visible over there, obviously, we will inform everyone at the same time through the exchange notification. So, we continue to see very strong demand, both from domestic as well as international customers.

**Shankar:** Just my second question, there's been a lot of talk about the GCC's off late and the surge in demand from that side overall across the IT industry itself. Are you seeing any trends with respect to the GCC itself in terms of the demand coming in from outside? That would be the second question. Thank you.

**Tarun Dua:** Okay. So, like, currently, of course, like, we have a few in the pipeline, but whether we can quantify that today, I don't know.

**Shankar:** Okay. All right. Thank you.

**Moderator:** Thank you. The next question is from the line of Keshav from Nivesh. Please go ahead.

**Keshav:** As mentioned on the website, there's one partner program. So, how much revenue do you generate from that? Like, what percentage of revenue did it help to, generate? You could answer that.

**Tarun Dua:** So, it's still a fairly small number compared to the overall revenue generated by doing. But then, like, we are working hard towards, like, building partnerships across the ecosystem. And hopefully, in the future, we'll see a lot more revenue coming out from this channel. So, it's more, I guess, future and forward looking rather than what we are seeing immediately as of today.

**Keshav:** And second one is, like, sorry, I joined a little late. So, you have already answered that. So, you have deployed 700 gigabytes in September and 256 in October. So, like, how much is the utilization rate? You have that number? You can share that.

**Tarun Dua:** So, I don't have it handy currently. So, but like, the last 256 CPUs, they were just are very, very recent. So, I would believe that, most of that would be unutilized from a revenue perspective, although they would be utilized for some purpose or the other. But, like, from a revenue perspective, I think most of that should be unutilized as of today.

And previously, whatever we have already have, like, apart from a few small outages where they are under some of the GPUs, machines are under repair and maintenance, I think, like, there is utilization anywhere between 80% to 90%.

**Keshav:** Thank you so much.

**Moderator:** Thank you. The next question is from the line of Kshitij Saraf from Tusk Investments. Please go ahead.

**Kshitij Saraf:** Hi. Good afternoon. Excited to see how the company is shaping up and excited to see the whole fund-raise go through and how it will be deployed. Question on the technology side, just out of curiosity. We have the Blackwell, which is coming along, and you mentioned you're looking to serve a new customer segment of source. So, with Blackwell, how will it work?

So, are you seeing economies of scale in serving the existing customers with this big GPU concept, or are you also looking to serve enterprise customers or geographical or more geographical diversification? How are you looking at it?

**Tarun Dua:** Let me try to create a path over here. See, if you look at the state of the art in the open-source LLMs, which are usually the heart of most Gen AI installations today, the open-source LLMs, they are the heart of the Gen AI installations. Now, if you look at the trend in the open-source LLMs, mostly we have seen that larger number of parameters in terms of number of billions.

So, a 400 billion, 405 billion parameter LLM in open-source would be very typical of the state of the art today. Now, that state of the art requires obviously bigger and bigger GPUs. So, today, the same LLM at 405 billion parameters would run at the minimum on a 16 GPU cluster, not less than that.

Now, there is already talk of 1 trillion parameter LLMs. And of course, there are people building the smaller footprint LLMs as well in open-source. It's not that the state of art for more specialized, verticalized LLMs is not there. But if you look at the major state of the art of the open-source LLMs, there is an increasing trend for more ebb and issue, zero-shot knowledge to be available within the LLM.

And they require bigger and bigger GPUs for sure. So, broadly, the trillion-parameter model that would become expand by the time the Blackwell gets released would be like whatever crystal gazing we can do, that seems to be where there would be a good match between the latest state of the art of LLMs and the Blackwell GPUs. And of course, if you look at the ecosystem, it continues to be very, very vast, where large number of organizations have large number of very, very different requirements from each other.

And people will continue to explore the entire open-source universe of what fits their problem space with a solution very, very well and continue to test and deploy it. So, I think ultimately, that's where the new demand for the Blackwells is going to majorly come from. So, very large LLMs in open-source, state-of-the-art, all the community editions released by the major LLM providers of today.

**Moderator:** Thank you. Ladies and gentlemen, this will be our last question. It's from the line of Richa from Equitymaster. Please go ahead.

**Richa:** Hi, my question has been answered. Thank you so much.

**Moderator:** Thank you. The next question is from the line of Karan from Keynote Capitals. Please go ahead.

**Karan:** Thank you for the opportunity. Sir, could you please clarify if the fund is used for GPU capacity or not?

**Tarun Dua:** See, broadly, I think like 75% of the funds will at least go towards augmenting the infrastructure capacity, which majorly includes the GPUs. And for the remaining 25%, we have taken that in the general corporate corpus, where there is sufficient optionality. Now, we would like to keep looking at opportunities that come our way for deployment of those part of the funds, some of which could go towards like helping increase our engineering and technology capabilities on the software side.

Other than that, like if we are not able to utilize for development of capabilities, or there are no other opportunities that kind of like arrive for utilization in the optional space, then we could like utilize that part for the infrastructure expansion itself.

**Karan:** Okay. And lastly, what are some primary internal KPIs that the management uses to monitor its performance, the company's performance?

**Tarun Dua:** So, we have also put that up in our latest presentation. So, we do intend to establish at least like one new major cloud zone, south of India. So, that would also include GPUs. So, we do intend to increase the diversity of our customer base. So, where we want to incrementally add more



government and enterprise customers. Obviously, scaling up of accelerated computing, GPU capacity and overall cloud infrastructure remains one of the key strategic initiatives for us.

And last but not the least, like we also intend to kind of like add to our engineering and technology and operational health in terms of both software and processes for delivering to our customers like the latest and greatest in the AI ML and AI technology.

**Moderator:** Thank you. Ladies and gentlemen, due to time constraint, this was the last question for today's conference call. I would now like to hand the conference over to the management for the closing comments.

**Tarun Dua:** Yes, thank you to everyone for participating in our call. And with this, we would like all of you a very happy Diwali. And thank you for the organizers of the call. Thank you, Megha. And thanks to all our customers and our team. And with this, once again, happy Diwali to all of you. And with this, we would like to end the call.

**Moderator:** On behalf of Go India Advisors, that concludes this conference. Thank you for joining us. And you may now disconnect your lines. Thank you.