E2E Networks Limited

CIN- L72900DL2009PLC341980

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

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

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

Date: July 24, 2024

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

Scrip Code/Symbol: E2E

Sub: Transcript of Analysts/Investor Earnings Conference Call for Q1 of FY 2024-25

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 Friday, July 19, 2024 at 11:30 A.M. (IST) to discuss performance of the Company for the 1st Quarter ended June 30, 2024.

The same shall also be made available on the website of the Company at https://www.e2enetworks.com .

This is for your information and records.

Yours faithfully,

For E2E Networks Limited

RONIT Digitally signed by RONIT Date: 2024.07.24 11:38:11 +05'30'

Ronit Gaba
Company Secretary & Compliance Officer
Membership No.: A59215

Encl: As above

E2E Networks Limited
Q1 FY25 Earnings Conference Call
July 19, 2024

Moderator: Ladies and gentlemen, good day and welcome to the Q1 and FY 2025 Earnings Conference Call of E2E Networks Limited hosted by Valorem 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 "*" then "0" on your touchtone phone.

I now hand the conference over to Ms. Purvangi Jain from Valorem Advisors. Thank you and over to you ma'am.

Purvangi Jain: Good morning everyone and a warm welcome to you all. My name is Purvangi Jain from Valorem Advisors.

We represent the Investor Relations of E2E Networks Limited. On behalf of the company, I would like to thank you all for participating in the company's Earnings Call for the First Quarter of the Financial Year 2025.

Before we begin a quick cautionary statement. Some of the statements made in today's concall may be forward-looking in nature. Such forward-looking statements are subject to risks and uncertainties which could cause actual results to differ from those anticipated. Such statements are based on management's belief as well assumptions made by and information currently available to the management. Audiences are cautioned not to place any undue reliance on these forward-looking statements in making any investment decisions. The purpose of today's earnings conference call is purely to educate and bring awareness about the Company's fundamental business and financial quarter under review.

Now, I would like to introduce you to the management participating with us in today's earnings call and hand it over to them for their opening remarks. We have with us Mr. Tarun Dua – Chairman and Managing Director and Mrs. Megha Raheja – Chief Financial Officer.

I now request Mr. Tarun Dua to start with his opening remarks. Thank you and over to you sir.

Tarun Dua: Thank you Purvangi and good morning, everyone. And thank you for joining us today. I would like to welcome all of you to the E2E Networks Q1 Earnings Call for the Financial Year ending

March ‘25. Let me briefly re-introduce E2E Networks and talk about the business of cloud GPU, especially for the first time listeners who might be new to the company.

So, E2E Networks was founded in 2009, and is amongst one of the largest Indian players in the cloud GPU market, which is based out of India. From our origin, like we have championed contract less computing in India since 2009. And we count a lot of current unicorns amongst our customers in their journey from startup stage to unicorn stage. Our infrastructure as a service platform supports a large variety of computing workloads for developers**,** SME, startup, enterprises, higher education and research and very recently, we have also been impaneled by MeitY, signifying the suitability of our cloud platform for use by the government for their compute workloads as well.

Our cloud platform supports a lot of cloud native services, including CPU and GPU based environment, running on top of virtual machines or containers, or sometimes in serverless mode. So we have a number of very flexible high performance storage solutions, including – object storage, block storage, parallel file system, container adapt storage, ability to elastic file system. We also support network based and advanced load balancing, virtual firewalls, NAT devices, queuing services, database as a service for a number of databases including like MariaDB, which is open source and Postgres SQL, and Vector databases.

We have continued to build our AI/ML platform called TIR, which is targeted at data scientists and AI/ML developers for AI/ML workloads, which includes training, inference, model end point deployment, building agentic framework, building RAG pipeline, and hopefully many more things to come in the future. So our entire platform is accessible online, via myaccount.e2enetworks.com via API services via as a terraform provider in the form of like command line tools that can be used by the developers and DevOps and data scientists to run any kind of workloads very similar to the way workloads run on any of the major hyper scalars today.

So as of today, majority of the workloads that we service now belong in the AI/ML domain, and majorly works on the cloud GPU infrastructure that we have built since 2018-19 onwards. Globally, cloud GPU has become like one of the most transformational technology opportunity for the coming next decade. And value creation which is envisaged by running AI/ML workloads on top of cloud GPU seems to be like at a global level being measured in trillions of dollars. Now, as a early mover into AI/ML, cloud GPU, and having built our own cloud platform from the scratch, we see a very substantial opportunity in front of E2E Networks. And we are very focused on delivering excellent value for money for our customers, while delivering the latest generation of cloud features that our customers have come to expect from cloud platform which includes basically the ability to run some 25, 30 different micro services. And a lot of back end delivery of services, which gets wrapped around in terms of being able to deliver these services with security and compliance features.

In the past one year or so there has been a talk of very substantial timeline, lasting over few months to more than a year for the delivery of hardware. Now, what we see today is that, the worst of the scarcity is now behind us. So we should be able to kind of plan our hardware run for much shorter period now, compared to what we did in the past, when there was a lot of scarcity. So, we continue to build our cloud platform features using a very strong software engineering team. And we hope to kind of continue to remain the first choice for India and many other countries hopefully outside India, for their cloud GPU workloads. And we kind of like continue building our platform based on the feedback on, what are the new features required, what are the gaps in the current features. And this productization where we continue to learn and continue to incorporate the features and bug fixes, is what generates the long term value and differentiation for us, as people get more and more used to our platform. And they have come to expect reliability, performance, and the kind of features that helps them kind of do the activities they are trying to do, run the workloads, they are trying to run, far more easily on our platform compared to other platforms. So that derives long term value for us.

We continue talking about the GPU driven computing. Now, how exactly is it different from the CPU driven computing? So in the CPU world what we see today is that, at a very reasonable price point, the number of physical CPU cores, which are delivered within a single server, are of the order of about say 144 or let say round it off to let's say 150. And typically you put two physical CPU dyes inside one server. So which means that we are talking about something 300 core. Now, compared to that, like a typical GPU server, it cost 10x, occupies maybe about 2x the amount of space, consumes about 10x the amount of power and depending on the kind of workload, so we will talk about the kind of workloads that GPUs primarily focus on which are majorly highly parallel workloads, compared to CPU which is working on more rule driven workloads. A GPU is typically able to deliver like, GPU has 17,000 into eight cores in a single CUDA cores or the GPU cores. So that's like, massively different from about 300 odd cores that a typical CPU driven server would have. And the GPU server also contains the CPU. So, essentially a single server for massively parallel workloads, driven by the state of the art in the AI/ML technique, majorly powered by open source is able to deliver anywhere close to 1000x, the kind of compute that take a typical CPU driven server is able to deliver.

So, which also means that, if you were to use the same number of cores or similar number of cores, the same order of the number of cores to try and run the massively parallel workloads, which are coming from the AI/ML domain, on to CPU driven system, you would have to actually do a lot of interconnects in terms of networking, you will need to set up like a large number of servers, occupy a lot of space. So, essentially what you are looking at is that your data center cost when running CPU alone kind of system goes up massively compared to running by very high density GPU compute.

Now, this is a trend that we also see in our own numbers, where we have seen that, from doing low density, CPU driven compute to high density CPU and then very high density compute on

GPUs, like our operational expenditures as a proportion to our revenue continue to come down. So where exactly are the GPUs used? So, typically the data that was kind of stored within organizational, within system was primarily driven by a rule based system, where essentially, what was really needed for those rules to work, shape the nature of the data. So, essentially, what you have typically in any enterprise was that, you would have relational data in the form of like a table, typically a single table would have say 500 columns. And then there would be entity relationships between this table and other tables, and the other tables would also have their own 500 columns, and then millions of rows or billions of rows. And then typically, what would happen in a typical enterprise would be that these 500 tables interconnected to each other, each with 500 columns or maybe at least 100 columns, then each of the columns denotes like very specific amount of data which could be typically be denoted through call, maybe just a bunch of numbers up to 10. And denotes like a flag, denotes like a very small piece of information, all connected together would constitute a data diagram for an enterprise software team, which would actually print it out on very large printers, cover up their wall of their Dev center and typically, people would be looking at that and writing a lot of software, rule driven software to kind of manipulate data going into this very complicated database driven system to follow certain set of rules that have evolved over organizations which have worked for decades altogether. So, the complexity would keep on adding, we would keep adding exceptions, we would keep adding a lot of rules to this system. And the CPUs were really up to the task of kind of like working in this environment and delivering like a lot of outcomes.

Now, all that is done with these rule driven systems essentially the organizations were really doing the kind of great job of automation. But on the other hand, if you look at most of organizations capabilities still belonged under people who would have their own set of mental rules, which could not really be translated into computer driven rules. So which is why we said that okay, if there is a computer driven rules but there are exceptions that needs to be handled, then you need a human in the loop. So, human in the loop was doing 90% of the work, automation was doing probably only 5% of the work and the organizations were working reasonably okay. Now, with GPU and AI/ML techniques becoming very, very modern, so what has happened is that, we have reached an inflection point of the amount of compute available for doing massively parallel operation, which can process a lot of probabilities in terms of like deriving inferences, very similar to the way in which the human brain works.

So, you can call these things like model, an artifact which essentially contain a lot of probabilistic rules, which kind of run on and an organization typically which was only storing 5% of the rule driven data can now actually afford to pass through this remaining 95% of the data through these GPU servers running state of the art, open source, software models or the AI model. And kind of start storing the knowledge derived from the passage of this data through these foundational models to kind of get fine tuned to deliver inferences, which are particular to one particular organization or one particular system in a way that it is more probabilistic not

rule driven. So which is like slightly more human driven and essentially, what we are saying is that, data itself is emulating the rules for these AI systems.

Now, what's the opportunity for India, we are talking about something really sophisticated, we have something, talking about leapfrog technologies. Now, India didn't have the wired broadband, like we leapfrogged into the 4G and the 5G world. And what that has resulted is that, as we are 20% of world's population, we generate 20% of world's data. Now, the opportunity in compute lies in the fact that, India only has 2% of the world compute capacity in terms of the amount of power, which has consumed by our compute and the amount of compute systems we have running in our data centers. So which means that, about 10x, the amount of compute still needs to be kind of designed and delivered for Indian organization, which has the capability to process almost 10x more data than we are currently processing.

So, if we are processing only like, so out of the 20% of world's data if we only have 2% of world's compute so there is a gap between the amount of data we have and the amount of compute we require to process it. So, which kind of presents the opportunity to leapfrog directly into the GPU segment, which requires a very different set of capabilities compared to what the CPU centric world required. So, this is a biggest opportunity that we are trying to pursue in terms of kind of helping India leapfrog from a GPU world to the CPU world.

Now, so what has happened, the whole world has obviously realized this, there are obviously developed countries, which move far ahead in advance of what happens in India. And the broad trend we have seen in terms of shift in the computing has been that earlier the cost of compute compared to the revenue of **(Muffled) 18:18**, typically was low enough, that people wanted its software capable of doing everything from everything, without worrying about the cost. So there was a broad trend of all the compute getting concentrated with very large hyper scalers, who charged a lot for their compute, because they were delivering the software capabilities of being able to service every single geography, every single workload from $5 to $0.5 billion per month and the R&D cost of doing every single workload in the world for every single industry, every single organization required a lot of customization. And most people ended up paying for the compute on hyper scalers, which was about 3x or 4x times of if they did the compute on their own and there was no space in the middle basically.

Now, what has happened is that, with GPU compute increasing the percentage terms of expense for any organization in terms of the amount of compute they used to deliver the software stack within an organization, that number is now drastically going up, which means that there is suddenly a space for specialized cloud providers who are not as expensive as hyper scalars. But in terms of capabilities, they provide almost a similar functionality, maybe doing 10% of the features, but delivering 90% of what is needed by most organizations. So, we are one such organization in India and some of the organization's in the US have now shown that it is possible to build very large scale multibillion dollar businesses in competition with the hyper scalers if you deliver certain set of services in a fashion, which will be very similar to

them, and kind of adds value in terms of not having to do 80 or 100 or 200 different activities that a typical IT department of a very large organization used to do, kind of there is a new space in the middle specifically for cloud GPU providers at every geography where people understand that what needs to be built, which is very similar to what the hyper scalers have built, but at half the price or lower, and since there is a massive amount of money going into the GPU compute compared to what was going into the CPU compute, where we were talking about cloud as a couple of $100 billion worth of business. Now, we are already talking about, GPU cloud as something like adding to the overall business of the cloud. So, which means that this has to come from the pockets of enterprises and organizations.

Now, the advantage of having specific cloud GPU kind of compute players is that, they can deliver this cost advantage very similar to the way they were delivering in the CPU world. Now, what was not proven in the market was that in the CPU world, we didn't see very huge examples of very large businesses built in a short span of time. Now in the cloud GPU world in the US, and in many other countries, we are already seeing specialized cloud GPU providers growing very big, very fast on the backs of this new demand coming in for GPU driven computing. So, that is the category in which your company operates. And now, I would like to hand over our call to our CFO, Megha who will briefly touch upon the financial and operational highlight of the quarter under review. Over to you Megha.

Megha Raheja: Thank you Tarun and good morning everyone. Let me first start by giving you some of the key financial highlights. I will summarize the performance for Q1 FY25. For Q1 FY25 the revenue from operations stood at around 417 million, which witnessed a substantial growth of 112% on year-on-year basis. EBITDA for the quarter is 274 million, which further provides a growth of around 168% on year-on-year basis. EBITDA margin for the quarter is 66.34%, which demonstrate a growth of 1403 bps year-on-year. PAT is reported at 101 million, which demonstrate growth of 44% year-on-year basis. PAT margin for June quarter is 24.46% and diluted EPS is 6.75 for the quarter, which is around 42% year-on-year increase.

If we do a quick comparison from last quarter that is March, we have witnessed revenue growth of around 40% on quarter-on-quarter basis, from 296 million to 417 million in the current quarter. EBITDA for the quarter is 274 million, which shows a growth of 79% on quarter-on-quarter basis. And we have reported a net profit of 101 million. That concludes the update for the quarter. And we can now open the floor for question-and-answer session.

Moderator: Thank you very much. We will now begin the question-and-answer session. The first question is from the line of Pratik Singh from DAM Capital. Please go ahead.

Pratik Singh: Sir, my first question is largely on unit economics on rental perspective. So, I understand that you will be paying rentals to your co-locating partners. So, do we pay rentals on a per rack basis or per kilowatt or megawatt basis and ballpark, how much would it be?

Tarun Dua: The individual figures are broadly confidential and covered under NDA typically, but roughly when we put everything together, then broadly kind of like that comes out to about, incrementally for new high density workloads, it comes down to about let's say, 10% or so of our front end revenue on a unit basis.

Pratik Singh: Understood sir. And the second question sir, as you rightly mentioned that the GPU based compute servers consume way more power and the power density for those racks used to be much higher. From what I understand is that the typical CPU base racks generally are like 10 kilowatts per rack max and GPU based was 50-60 kilowatts per rack. So, are our co-location partners do they have this kind of an arrangement where you can fully fit your 8x of the GPU servers on one rack, or it's only a traditional arrangement where on one rack instead of filling it completely, you are just filling one 8x extended server which takes around 10 kilowatts. So the question that I'm asking is, because we didn't have a lot of AI based centers in india. So, how is it working right now, are we not completely billing the racks because then that will need a lot of cooling and all, so how is it working right now?

Tarun Dua: So, I didn’t really understand the question, but like, I'll try to answer it as best as I can. See, most of the data center facilities including like any facility, which was planned to be built even like say four years back, would not be built for very. Anyway, was not very useful in very warm climate of India. So it used to make a lot of sense to do density of no higher than about 14 to 15 kilowatt per rack. Now, does that mean there is miss utilization of space, I probably wouldn't characterize it in that way. Because you are not just running GPU servers, you are also running CPU servers, you are also running testing fabric, you are also running storage, we are running a lot of equipment. So overall, any data center has to be designed with generic workloads in mind. Now, is it still able to run the GPU servers, the answer is yes. Like most of the data centers, today at least have ability to run the GPU servers. Without any overall significant difference in the cost compared to a very high density data center.

Pratik Singh: Understood sir. In the cooling requirements and all, the cooling facilities don't need to be changed, they are as it is that they were earlier.

Tarun Dua: Not for another one or two generation. High density compute, off the kind we are looking at they are still about a year and a half to two years away.

Pratik Singh: Great, understood. And sir my last question is largely on H200, which is now being launched? Would we be going to them, or we think that H100 for now is enough sufficient for most workloads that we are seeing right now?

Tarun Dua: Okay. So like you have to match the workload with the price per unit of teraflops. So, in that sense that typically X100 could be slightly more expensive than H100. Now, depending on the workload, if you do not have typically more memory requirements, you will continue to be able

to use like H100, even for new workload without having to spend extra for the extra amount of GPU memory.

Moderator: Thank you. The next question is from the line of Prathamesh Tiwar from Tiger Assets. Please go ahead.

Prathamesh Tiwar: Sir, I have three questions firstly, on the CAPEX side. So, we have announced around 800 crores of CAPEX for FY25.

Tarun Dua: No, we didn't actually that is not true. We didn't announce 800 crores of CAPEX, and somebody asked, how much would you like to spend, so my answer was that, in an ideal world if we had the money, you would like to spend that much. So there was no announcement from our side as such that we will do 800 crores of CAPEX. It might end up that we end up doing that much, but let's not call it an announcement. So anyway, so that was the first question I guess.

Prathamesh Tiwar: Yes, so sir let's say you are going to do but some amount of CAPEX is expected right sir for this year. So just wanted to know how it's going to be funded?

Tarun Dua: See broadly it's going to be a mix of vendor financing debt and internal accruals and potentially some amount of equity raise.

Prathamesh Tiwar: Okay. So sir any figure you would like to give on the CAPEX side for FY25, looking at the economic scenario?

Tarun Dua: Sorry, I didn't understand.

Prathamesh Tiwar: You said 800 crore CAPEX is not fixed as of now, so any figure you would like to pick?

Tarun Dua: It could be more it could be less, so we are very, very flexible. So this is no longer the environment where you had to wait for hardware for six months so you had to really plan ahead. We used to do a lot of just in time in the past and the planning cycle for us, we can drastically shorten it to about six to eight weeks as opposed to having six months or 12 months kind of planning cycle now.

Prathamesh Tiwar: Okay. And sir second question is on the MRR side. So, you have told that you will do MRR of around 14 to 16 crores for FY25, if I'm not wrong. So if we see in Q1 we have done MRR of around 13.5 crore, so how are we looking at coming quarters, will coming quarters be more towards 15, 16 or what number you would like to give?

Tarun Dua: I wouldn't like to give any number, we are never in a habit of giving any guidance for next quarter or next-to-next quarter, so we are talking about next three quarters. So we will simply have to wait and see. Also, I would like to correct you that the MRR for June is not 13.5, the MRR for June is about 14.5.

Moderator: Thank you. The next question is from that Pratik Chaudhary from Samarth Capital. Please go ahead.

Pratik Chaudhary: Sir, for the as you said that zone MRR is roughly 14.5 crore does that largely mean full utilisation of the 450 AI H100 units that you have?

Tarun Dua: I would say about let say, about 90%, 95% utilization.

Pratik Chaudhary: On MRR?

Tarun Dua: So, we still have additional inventory available. So, to the extent that this is about, at about 90%, 95% of the utilization.

Pratik Chaudhary: In your rough assessment how much, what is the time in terms of number of H100s or similar such devices required, what was the India requirements coming along in the next one or two years?

Tarun Dua: These are all wild guesses ultimately, because in a sense that, when you start seeing the industry reports by that time all estimates, are based on the actual numbers. So, India in a way follows the developed market. So, we are very, very hopeful that what we are seeing in the developed markets India will also see that and especially too because the gap in India in terms of processing capability is far higher. Like you can't predict what will happen in a quarter or what will happen in two years, but we are very, very hopeful of a very bright future for cloud GPU industry in India, based on the fact that what we are seeing in the developed markets today.

Pratik Chaudhary: Are we also in a position to target the international market, is that already happening from the current GPU?

Tarun Dua: Yes, that is already happening. So in the sense that essentially this is the business of solving problems. So like we are already seeing some level of traction outside India as well.

Pratik Chaudhary: Roughly what percentage of your AI diffused would be deployed for overseas customers?

Tarun Dua: So it would be premature to comment based on say one quarter or so, but if we could predict it, it would be like over a year or would be somewhere between say 25% to 30%.

Pratik Chaudhary: That’s a great number and what’s the pricing difference as of now vis-à-vis your international competitors for the same DPU AI?

Tarun Dua: Indian customers and international customers pay the same.

Pratik Chaudhary: No, the price that your competitors offer, your international competitors, what is our roughly percentage?

Tarun Dua: There is never 1:1 price comparison, because there is always tremendous difference between features of one provider versus the other. So there would always be provider who charge like almost like 5x of what we charge, there would always be providers who charge like maybe even 25%, 30% less than what we charge. So we are like never the lowest priced compute provider out there because we have a whole set of features that we have built and capabilities that we have built. Also, the pricing depends a lot on a lot of things in the way in which a customer typically consumes the compute process. So, it's not like, it's never an apples-to-apples comparison when you say that, one GPU versus one GPU, what are you paying X provider versus provider Y?

Pratik Chaudhary: Last two questions. There was an announcement that B100 would be coming in sometime and the difference in terms of compute, without it it's almost 2.5x better with only 25% increase in cost?

Tarun Dua: Again when you talk about software, these are very, very absolute numbers. So, the difference would be 25x for modern software built with modern techniques of that era, in the future, the difference could be as low as 1.3x for all you know and care. So, you cannot kind of place a theoretical number and say that okay, this is the projected theoretical number. So based on this theoretical number, this is what, by what number the compute has advanced. So, a lot depends on the software stack, which is being run. And the ability of the customer to kind of derive the benefit also depends on how much they are willing to invest into re-architecting their software to make use of the new hardware that gets delivered in the future. So it's never cut and write that way.

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

Kshitij Saraf: Just zooming out a little bit, how is the market shaping up in terms of you making the buzz around providing the GPU and obviously the competition, at the overall level when we look at it is increasing yes, but at the same time you mentioned there is no one provider and there is no vanilla solution as such. So just some color on how you are differentiating, or what are the use cases that you are building as we go along?

Tarun Dua: We are very focused on national language processing and computer vision and their equivalents in the generative AI world which is the LLMs and diffusion based model. And apart from that, of course we will continue to fill the gap in supporting like newer workloads. Now, essentially the way to look at this business is not from the lens of asset monetization, but from a software platform perspective. So as long as you keep building a software platform which serves the needs of the customer, the customers will continue to use the platform. So that's the whole idea behind the business.

Kshitij Saraf: Okay, understood. So is this software part only, how are we trying to sort of stay ahead, is there a plan that you guys have in place or because the situation is rapidly changing, you are just trying to make sure that you have the software updated in terms of the updated use cases, the requirements that are coming up?

Tarun Dua: So, we obviously have a very active development team, which continues to release features like every couple of weeks. And every quarter we make a lot of progress in terms of what we are doing for our entire cohort of customers, based on the learning's that we get from our interaction, both from industry experts, our customers, and what is the state of the art, ongoing in the AI/ML community in general.

Moderator: Thank you. The next follow up question is from the line of Pratik Chaudhary from Samarth Capital. Please go ahead.

Pratik Chaudhary: Sir, were there any one offs in our quarter one?

Tarun Dua: Sorry, I didn't get your question?

Pratik Chaudhary: Were there any expenses that we capitalized in Q1 FY25?

Tarun Dua: I will let Megha answer this question it seems like a finance question.

Megha Raheja: No, there are no expenses which have been capitalized.

Moderator: Thank you. The next question is from the line of Abhishek from ABC Capital. Please go ahead.

Abhishek: My first question is, you had given guidance of approximately 40% revenue growth year-on-year for the medium term. So in this year, will you be doing any upward guidance?

Tarun Dua: Like I said, we don't usually give any guidance. So broadly, what happens is that, that is usually the interpretation of the people asking the question. So, in an ideal world you don't know what is going to happen in the next quarter, and next-to-next quarter so, the cloud business obviously is not a business of consistency at very small scale. So, ultimately, we will have to always look at the past and predict the future.

Abhishek: Okay, thanks. Another question is, are the promoters thinking of any, of selling any stock in the near future?

Tarun Dua: Again, you are asking me to predict future, but broadly whatever we do will obviously inform the exchange and inform everyone.

Moderator: Thank you. The next question is from the line of Aditya Trivadi from Nippon Capital. Please go ahead.

Aditya Trivadi: One question I had was, was that approximately your ARPU quarter-on-quarter has gone up as much as 48%. So, if you could just give us a breakdown of what you attribute that to as well as, what the client addition has looked like for the quarter?

Tarun Dua: Basically, from a customer perspective we have the kind of like, although we do not really differentiate as a very small company between the customers but, what we have seen is that the larger customers have grown for us in terms of where we have added the revenue. And this is primarily driven by the cloud GPU side of workloads for typically like AI/ML kind of like workloads.

Aditya Trivadi: And in terms of client addition, how many more clients added this quarter in terms of new clients?

Tarun Dua: So new cloud GPU world, we don't really count the customer addition, but the amount of workloads we have added, because we have moved away from the era of trying to collect a lot of small customers. So we are moving towards kind of higher ARPU customers. So, broadly in terms of your numbers, I don't think we have seen any significant addition in terms of the overall expansion of the overall number of customers that we have, but what we have seen is that the new logos on the top end of the revenue that we have increased.

Aditya Trivadi: Also in terms of customer concentration and revenue, what percentage of your top 10 customers will be attribute to the current revenues?

Tarun Dua: See, this is again like a dynamic number, in this quarter probably the concentration would be higher than usual, so top 10 customers might have gone up by all the way compared to previous quarters to say 45%, 50%. On the other hand, like this is a number that should be looked at in retrospective at the end of the year rather than on a single quarter business. Because, the cloud GPU workloads are typically like bursty workloads, where a single customer might come up with a short burst of large amount of compute usage. And next quarter that customer might be kind of significantly lower than this quarter.

Moderator: Thank you. The next question is from the line of Rajshekhar an Individual Investor. Please go ahead.

Rajshekhar: So, I wanted to ask about the CAPEX spend for this quarter, it seems to be about 23 crores. So is it a reflection of the demand that you are seeing, as you said that you plan for around six to eight weeks lead time is required. So the number seems a little lower than the last couple of quarters, that was my question.

Tarun Dua: So again, it's kind of like you are also looking at the tail end of having done a massive expense in the Q4 of the last year. So, in that sense we built a lot of capacity during that era. So

obviously, we will not be consistent every quarter. So, it depends on local minimas and local maximas in many ways.

Rajshekhar: And what do you think, do you have an estimate for this quarter at least, not for the entire year, but for this quarter?

Tarun Dua: This quarter, next quarter, those kinds of estimates we do not give any sort of guidance on call.

Moderator: Thank you. The next question is from the line of Vishan Joshi an Individual Investor. Please go ahead.

Vishan Joshi: Sir, I have two queries. First of all, are we participating any bidding with the government, because as you have been enrolled with the MeitY

Tarun Dua: Currently we are not participating in any bidding with the government. So we look at the government business in three different ways. One is where we want to support our MSME partners who typically kind of generate business from the GEM portal. So on a case by case basis, we want to work with those MSME partners, it's only been less than a month since we got certified by MeitY. So, it's very, very early to say what sort of business we will be able to generate over there. Now, the second part is where we will start evaluating where there is a very good match between our **(Muffled) 48:49** and those kind of workloads, when they come about in form of like requirements from the government, where we have a good match in terms of size and capability. So that is something that will happen over next, I would say 18 to 24 months. So, we are looking at a medium term over there. And the third possibility is that where our capabilities are usable in conjunction with someone else, who has other capabilities, which they can combine with our capabilities and then we can jointly go with those kind of partners. So again, all these three things are far enough in the future for me to comment on them today.

Vishan Joshi: Okay, sir another query. If we see the financial results for this quarter, again the profitability is improved if we also see the operating expenses that is also dipped by a greater margin. So, means it's an aberration or it will continue to be in the same line?

Tarun Dua: Like we said that this is a broad trend because of the density of compute going up where the percentage of data center cost as a new will hopefully continue to do go down in the future as well.

Moderator: Thank you. The next follow up question is from the line of Aditya Trivadi from Nippon Capital. Please go ahead.

Aditya Trivadi: Just in terms of CAPEX, you are obviously not giving any forward guidance in terms of the amounts but if you could just give us the mix in terms of equity and debt that you all should follow going forward in terms of your CAPEX?

Tarun Dua: Any kind of equity raise we will obviously inform the exchange before we kind of pursue that. And of course, we will continue to use a judicious mix of internal accruals, equity, debt, vendor financing all of them in various ways to kind of finance our CAPEX.

Moderator: Thank you. The next question is from the line of Pratik Chaudhary from Samarth Capital. Please go ahead.

Pratik Chaudhary: Sir, FY25 what is the bare minimum CAPEX that would be achieved based on whatever assessment you have and what is the current terms?

Tarun Dua: So, that's again the same question. Broadly asset terms wise somewhere between 0.5 to 0.6. And it also depends on, what sort of SKUs the customer choose to deploy. So, there is a lot of variation in terms of not being able to kind of very accurately predict what is asset terms. Now, in terms of bare minimum, we don't have a kind of minimum number in our mind that okay, this should be at least done, it is always based on the current demand in the pipeline, based on which we will continue to make investments into the CAPEX.

Pratik Chaudhary: The asset turn has fallen from 0.7 to 0.5 to 0.6?

Tarun Dua: Somewhere between 0.5 to 0.6.

Moderator: Thank you. The next question is from the line of Richa from Equity Master. Please go ahead.

Richa: I just wanted to understand, beyond the hyper scalers what are the entry barriers, and what would stop more players from coming into this space or do you think that the opportunity is big enough that, the pie is big enough for everybody to grow?

Tarun Dua: Yes, definitely the opportunity and the pie is big enough for a lot of players to exist in the marketplace between various levels of value addition like various players would do and the main kind of like, I would not say a barrier to entry, but the main way in which people, in which the competitors for, the players in a particular market differentiate in this particular market is through the software. Now, in our case, we have been kind of been writing our cloud software for a good part of the last decade. And on the cloud GPU itself, we have been working for last almost four, five years. So we believe that gives us the head start in terms of, understanding the customer needs. And having help customers with a lot of problems, it gives us a leg up in terms of understanding, how to kind of build around it.

Richa: Sir would it mean a lot of client servicing or engaging with client or in that sense, would it require some kind of?

Tarun Dua: This is a product business, so doesn't mean a lot of client servicing. So this is ultimately not a services business, which is dependent on the number of people but it's a product business, which depends and relies on the capabilities of an engineering team, which works in tandem.

So where just trying to understand from customer that, okay what is it that they are trying to achieve. And obviously, one part is to build a solution based on the current technology and the solutions that we have, and then continue to help the business, identify all those gaps, and then fill up those gaps in terms of like, build out of new features and new software, and kind of like new integration.

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

Kshitij Saraf: So, Tarun just to get a sense of your overall vision for E2E now, obviously we have the GPU offerings, and we have the software's around. So do you plan to keep evolving the software's around the increase hardware requirements that come up is that something when we think of E2E five years, seven years down the line is this something which you would be continue doing or do you have sort of another?

Tarun Dua: Product development is like a continuous process, because the cloud world, and especially like a rapidly evolving field, like AI/ML it changes and constantly evolves. So which means that the software that we do, will have to constantly evolve. So that is something that is not going to change. And then, obviously to kind of solve at adjacency of the problems that we are solving today. Like we kind of continue to build more features. So for example, let's say there are X number of customers in a market and today our software is able to service let's say Y percentage of those, obviously you keep pushing the envelope from Y to X. And also, to figure out in adjacent countries or different geographies or any other adjacency that whether how you can easily expand into another market by building new software. We are big believers in that and we will continue to do that.

Kshitij Saraf: Okay. And in terms of what is the specific software's are you planning to go deeper on specific verticals, how is the plan there?

Tarun Dua: So, obviously there are about like 40, 50 different paths, so the number of workloads that can be using AI/ML, in different ways and the foundational technologies needed for those is a very large set of software. So we will continue to kind of evolve our software game into all the variants.

Moderator: Thank you. The next question is from the line of Hardik Satya an Individual Investor. Please go ahead.

Hardik Satya: So my question is related to the CAPEX that we did last year of 185 crores approximately, what is the status of deployment, have we done the deployment completely up to April or was it done up to March or up to May so that we can get some MRR clarity?

Tarun Dua: See broadly, the entire deployment was done by somewhere around end of March or maybe first week of April. So, almost all of that is done, deployment means that it is available on our cloud for the customers to be able to see and in terms of MRR conversion, the overall inventory that we acquired including this as of today, from that inventory we have utilized about let say, including all the other inventory, we have utilized overall about +90% of the inventory.

Hardik Satya: So given a 90% utilization and in the near future do you see?

Tarun Dua: We have also mentioned in our presentation that we are already in the process of acquiring another 256 H100 which would go live in the quarter.

Hardik Satya: So given the forecast for the near term, if money or funding was not a problem or a challenge, what would ideally be the demand that you are foreseeing. So would it be to the quantum of 800,000 or maybe 2000 crore also?

Tarun Dua: No, so that results in people misquoting me, a formal announcement that this guy.

Hardik Satya: No, I would not take it as an announcement.

Tarun Dua: I will not take that bait this time, and I will say that let's wait and watch. We will keep everyone informed of our plans as they get formalized.

Hardik Satya: But I can take that much that you are seeing a good future ahead and there is a significant demand is what I can take from your reply?

Tarun Dua: Yes, so over let say medium term, long term, over like next decade, like GPU driven computing is the new thing, which is like very transformational for entire industries and potentially for seeing new industries that we have not even seen till now. So we are very gung ho on AI/ML and GPU driven computing world.

Hardik Satya: And also on the fact that government is pushing for having the data centers in India and keeping the data within the boundaries of India, what kind of opportunity do you see in those segment?

Tarun Dua: That's an evolving trend. So that helps us but it's not quantifiable that to what extent that helps us. But we do see some effects of that, we are now seeing a trend of like more and more Indian companies recognizing the need to work with Indian vendors. So whether it becomes like a huge snowball, we don't know as of yet, but we are seeing green shoots of that behavior, that as an Indian company we are more comfortable dealing with an Indian company rather than company from outside India.

Moderator: Thank you. The next question is from the line of Aman an Individual Investor. Please go ahead.

Aman: My question is, what is our USP over the, if we are in a competition with the largest player or an emerging domestic player, somewhere what we see the competition that makes us different from the competitor, which is infrastructure heavy player and what add on's we provide in terms of cloud services to our clients?

Tarun Dua: So there is a long list of cloud services that we have been providing from last many years. And typically, what happens with software usage is that software usage is a habits, so whether it's a larger player with like a bigger piece of software, whether it is like a player which is more infra heavy, less software, ultimately the differentiator would become the stickiness of the software that we eventually produced, with the help of like understanding from our customers, that we build from our customers. So ultimately the differentiator is going to be the software platform that we are building.

Aman: So are you trying to convince me that we are.

Tarun Dua: I am not trying to convince you of anything.

Aman: No, are we at better.

Tarun Dua: I only want to convince our customers to use our product, I don't need to convince anyone else.

Aman: No, I just wanted to tell that, what I am understanding that we are at a better stage in cloud services software, from what you are telling?

Tarun Dua: We initially launched our very first version of the software, way back almost eight to 10 years back. So obviously, our software has seen like 1000s of customers running their production on that. And, of course we have faced a lot of issues and feature requests, which you have called for. So that gives a software like definitely a lot more maturity than many of our competitors.

Moderator: Thank you. The next question is from the line of Gurubaksh Singh an Individual Investors. Please go ahead.

Gurubaksh Singh: My question is, you mentioned to one of the previous queries that, we have been moving steadily towards the larger customer profile. And as we do that, do you think that we would be taking on the bigger players in the competitor set and would that impact our margins going forward?

Tarun Dua: See, enterprise software and enterprise cloud is a very permission driven market. So as you solve deeper problems, your margins ideally like, get somewhat impacted by the amount of commissioning you do initially. But eventually, it results in much higher ARPU than much higher margins. Because the amount of stickiness that we produce in terms of lifetime value of the customer continues to increase.

Gurubaksh Singh: Okay. So are we saying that as we improve on our solution based approach, there is a possibility to even get higher margins, is that even possible?

Tarun Dua: Yes, definitely.

Moderator: Thank you. The next question is from the line of Pankaj Kumar, an Individual Investor. Please go ahead.

Pankaj Kumar: So, only one question I have, so the margins have really improved the EBITDA margin, is it something that is sustainable in future quarters or in upcoming years?

Tarun Dua: Ours is a platform business so, the overall team size, the platform size doesn't need to really increase with the amount of cloud business that we do. So, we can potentially scale up the business 10x where the platform remains practically the same in terms of the size of the dev team and the size of the operations team, it needs to grow very organically compared to, whatever the growth on the side of the content revenue. So that way, there is always any platform business, the scope for growth of margins is always there.

Pankaj Kumar: Okay. So, similar margins can sustain in the future quarters also for FY25, is that understanding correct?

Tarun Dua: I don’t say quarter-by-quarter I would say medium term and long term broad trend should sustain.

Pankaj Kumar: Okay. So, can we say that we have moved from +50 percent kind of margin to +60 percent kind of margins, is that a reasonable estimate or reasonable understanding?

Tarun Dua: I already answered this question, in the medium term and the long term the trend should sustain.

Moderator: Thank you. We will take this as a last question. I now hand the conference over to the management for closing comments.

Tarun Dua: Thank you, everyone. I always continue to learn from all the questions that we get from our esteemed investors. Thank you, all of you, thank you to the Valorem team. Thank you, Megha and everyone for joining us today. Let's hope to continue our conversation over the future calls. Thank you, everyone.

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

****