

How Lenovo's CFO Is Allocating Capital During One of History's Biggest Booms | Odd Lots

56 MIN • YOUTUBE • [HTTPS://WWW.YOUTUBE.COM/WATCH?V=KXYEEK8IXPM](https://www.youtube.com/watch?v=kXYEEK8IXPM)

<https://www.youtube.com/watch?v=kxYeeK8IxpM>

SUMMARY

The Odd Lots podcast features a conversation with Winston Shang, CFO of Lenovo, discussing the company's role in the AI infrastructure landscape and the implications of token usage in corporate budgeting. The discussion covers the integration of AI in hardware, the importance of supply chain management, and the evolving nature of capital allocation in tech firms.

- The term "token" has been added to the Chinese Mandarin dictionary, reflecting its growing significance in both monetary and linguistic contexts.*
- Lenovo positions itself as a global AI infrastructure provider, emphasizing its ability to deliver AI solutions from device to cloud.*
- CFOs are increasingly focused on managing operational expenses related to AI token usage, with a need for clear budgeting and return on investment metrics.*
- The competitive landscape in AI is marked by a divergence between US and Chinese companies, with Chinese firms often benefiting from lower costs due to intense competition.*
- Lenovo aims to optimize its supply chain and manufacturing capabilities to meet the growing demand for AI infrastructure, including data centers.*
- The podcast highlights the importance of flexibility in AI partnerships, as companies seek to leverage various AI models without being tied to a single provider.*
- The discussion touches on the challenges of securing components for AI infrastructure, noting ongoing supply chain constraints and the need for efficient energy sources.*
- Lenovo's strategy includes balancing shareholder returns with investments in innovation and growth, particularly in the AI sector.*

[music] >> Bloomberg Audio Studios. Podcasts, radio, news. >> [music] [music] >> Hello and welcome to another episode of the Odd Lots podcast. I'm Tracy Alloway. >> And I'm Joe Weisenthal. >> So Joe, we are here in Hong Kong still. And we're learning a lot of different things, a lot of interesting things. One of the most interesting things we heard actually came yesterday from the Baidu CFO, where we were just talking casually ahead of our interview, but he was saying that the word token >> Yeah.

>> has now officially been added to the Chinese Mandarin dictionary. >> And that the characters that actually make up the Chinese word for token are something like word currency. >> I think that's so fascinating. You know, I'm fascinated by the etymology of the word token specifically. So I was thrilled to hear that. But it is, you know, it's like we know the word token in the monetary context. We know Chucky Cheese tokens. We know [laughter] crypto tokens. But we also know that since the middle of the 20th century, linguists have been using

token to describe more or less a word. And then obviously with LLMs, we talk about these linguistic tokens a lot. So to see that in Chinese, the formal term is a merger of these two concepts is I find a very intellectually satisfying thing to learn about.

>> Bringing up Chucky Cheese tokens is really a way to make sure everyone knows that you're a millennial, Joe.
>> Yeah, that's right. That's right. Arcade I I should say arcade token. >> That's right. >> Yeah. >> Okay. And I think what's really interesting about word currency itself is it implies that it's connected to spending, right? >> And so when you think about the big companies that are spending all this money on tokens and I guess the AI infrastructure build out more broadly, CFOs become really important, right?

>> Totally. So, we, you know, the headlines by and large are about uh CapEx, right? And that's going to continue to be because of the data center build out. We're going to talk about that. But a lot of it is also going to be OpEx and figuring out how to within a company allocating token permissions and caps and so forth and I doubt anyone has figured out the final answer. But if two different people can get different value out of using AI models, then there is no way that it makes sense for them to have the same token budgets.

By the way, Tracy, can I ask you a personal question that I've never asked you before? >> Oh, okay. >> Outside of the work context, are you a Mac or a PC person? >> Oh, I mhm, I only have I only have work computers at the moment. >> Okay. All right, before that, but >> Before that, before that, definitely PC. >> Okay, good. Me, too. >> Yeah. >> And in fact, >> it's not it's not a choice, is it? Like, I Sorry, I I really don't like Macs.

>> I don't, either. And in particular, I am long been a fan of the what used to be the IBM ThinkPad laptop with the famous red button, which is now owned by Lenovo. >> That's right. You've talked about that computer before, so I can say that this is categorically true. Joe likes that computer. But what's interesting about Lenovo is like, okay, it's famous for the computer with the little red button in the middle, but it's now making an AI play. I mean, everyone's making an AI play, but it's doing it from a different perspective. So, AI integrating into the actual computer, the hardware, but it's also doing cloud, right? So, this is a really good opportunity to, I guess, take the temperature on AI, the AI build out, the AI spend from a bunch of different perspectives. So, we do in fact have the perfect guest.

We're going to be speaking with Winston Shang. He is the CFO of Lenovo. So, Winston, thanks so much for coming on Odd Lots. >> Hey, thank you, Tracy, and thank you, Joe. >> Why don't you go ahead and describe what Lenovo is and what the AI play is and how it actually fits into the existing business? >> Yeah, Lenovo is a global AI infrastructure company that provides pocket to cloud AI infrastructure for the consumer and the enterprise. I think that's really in a nutshell. And so, today we're able to provide this to hyperscalers, which are doing a lot of the spending, are driven by training demand today. And then, given our IBM x86 heritage, which we also acquired the server from the IBM actually. So, we actually are very strong in CPU compute as well, given that IBM was a dominant player in the x86 architecture.

So, from that perspective, we're well positioned for inferencing needs of the enterprise and also the hyperscalers in terms of in the cloud as well. So, I think from the perspective of then, you talked about tokens. And I think

from a token perspective, really today people are in the early stage of how much I'm really paying for. I heard someone saying that there was an engineer at a particular company, which I will not mention, that apparently spent \$100 million in a month on tokens.

And so, as a CFO, I would have concerns cuz that was clearly not in the budget, right? Not saying that that was from Lenovo. So, we cannot have that happen. And I think we need to be able to drive the productivity or efficiencies as it relates to that budgeting of the token generation. And I think a lot of that will happen on device, where you may just pay a higher price for a device, but you know what you're going to be able to do on compute for the security of the data and for the privacy that you want to interact with your AI agent.

>> So, just to be clear, for the computers themselves, they can do some inference, right? But where it makes sense, you route it to the cloud. Is that right? >> That that is our goal. So, the Lenovo Agent AI today we what we are good at is really integrating and maximizing the compute capabilities on a device. And from that perspective, given the various needs in terms of various operating systems, agents that may not want to sit on top of a device, I think our agent aims to orchestrate the various LLMs. We'll take the compressed versions of these LLMs.

We will, depending on the partnership, be able to do the on-device compressed LLMs versions and do that as a local compute, but in certain queries, allow it to go on the cloud. And therefore, probably would allow the user probably to spend in terms of the token generation. >> It's interesting this word orchestrate because you so you have the AI agent orchestrating maybe a bunch of different sub-agents to complete some task, and that is something that perhaps is done best on a CPU. A company that builds servers is also an orchestrator of the supply chain and acquiring the different components that go into a server, etc. I want to go back to something you said in your first answer because I think this actually will get to the core of this new era. You said, "Okay, an engineer spends \$100 million in a month on tokens." And it's like, that would not make you happy as a CFO. But it could make you happy, right? What if you had a multi-year database migration plan that you think, "Oh, this would be a \$500 million job," and the engineer does it in a month for \$100 million via tokens? Don't you have to at least be open to the possibility that that was money well spent?

>> Absolutely, Joe. I think everything is about the return, right? And the planning. So, we're not afraid to invest. As a CFO, you are there to allocate capital. You're not there to constrain capital. You have to allocate, but you have to be clear in terms of that return. And I think in that case, it's probably one where they weren't sure in terms of that the that what they were doing in terms of that spending and wasn't in budgeting. And that that goes to the point of what is happening today.

I think most enterprises were at the early stages of how people are changing from a subscription-based model to a token usage model in terms of the compute capabilities. And so, that is at the beginning and enterprises are starting to figure out, how do I really track that spend? And therefore, what do I really want to get out of that return? And so, we're on that early stages. So, I you're absolutely right in terms of that we would want to spend if we can get a return out of it. You're absolutely I absolutely right. And I think preliminary data and anecdotal evidence from a lot of people in that I chat recognize the power of AI and the probability of

generating increased productivity and potentially even cost savings.

>> Wait, say more about measuring the actual return. And again, I realize it's early stages, but I'm curious how you think it will actually be done because we've been asking a bunch of executives, how do you internally benchmark your AI projects or your token spend? And everyone says, "Oh, productivity or cost savings or whatever." But those are They sound very theoretical. So, is there more concrete, I guess, KPIs that you're looking at? >> Yeah, I I think it really depends on your specific enterprise.

In our case, I would really hone in on a few areas. For me in terms of pricing and that visibility of demand in terms of channel inventory and the demand of 180 markets and what I would take to actually provide and work with my channel partners in terms of managing that inventory supply. There's a lot of economics involved in that. So if we can optimize that to run by AI rather than spread out in 180 markets, I think there could be a lot of productivity gains from that that could be seen from a dollar perspective rather than just generally in terms of making very blanket general statements. Other areas I would say that you could also increase productivity is pretty clear in terms of data spend that you have today.

A lot of hedge fund friends and investor friends really talk about in terms of what their savings could be in terms of how much they're paying for certain data subscriptions. And then of course as it relates to IR functions, to M&A functions in terms of how you increase productivity. Those functions tend to be smaller in head count so it's really less about the head count reduction but probably about the increased productivity and therefore are you making the right decision because those are small head count high impact areas, right? You could have a higher stock price and therefore your cost of capital becomes much more efficient in your financing.

Even M&A if you have better analysis, you're optimizing in terms of how you're returning in terms of that capital deployment. So from those perspectives, I think it works out. And then of course areas like marketing where we're spending a lot out there in terms of you know, video and other production generation. Those actually can be done by AI today. So we really need to look at that. So there that's just a few examples but I think it there must be hundreds and hundreds and we're in the process of [music] trying to sort through that.

>> [music] >> Let's go back to the theoretical engineer who spends a hundred million dollars on tokens in a month because you know, it seems to me like you're going to have someone who actually can with their you know, intuitions and skills with AI etc. Actually deliver ROI even if the nominal amount they spend is nosebleed and then other people are just going to waste a lot of money and I guarantee I know a lot about wasting money on AI and I guarantee that it's very easy to develop that psychosis whereby you believe you're doing some incredible things that feel like magic and productivity and it turns out that you really haven't built anything etc. Like it's very easy to delude yourself into thinking you're being productive with one's AI use. What I'm curious about specifically is how do you even go about say identifying the employees or the teams who might really be people who you really should give a long leash to with their spend. How do you even like identify who those are that should uh yes, let's not be too aggressive in their consumption in capping their consumption.

>> Don't give all your tokens to a Joe with AI psychosis. >> How do you avoid giving all of them all giving me

the uncapped budget and finding that person who actually deserves a very liberal spending budget? >> Yeah. And you know, these are great questions. They're so relevant particularly in a large enterprise like us in a global Yeah. for global 200 company, right? In 180 markets, 80 plus billion in terms of revenues. We have massive scales and so many tens and thousands of employees. So, it's really difficult from a CFO's perspective to micromanage down to the person's allocation. So, you really have to be able to say, first of all, what are the agents in each domain capable of doing today, right? And how what is that productivity and also capability?

And then you can say, well, what specifically can it do today? And then let's get the dollars, at least from the CFO perspective, let's then eliminate the dollar spend in those areas and really force the AI. I have come to the probably the conclusion that to really effectively drive things, I need to then force discipline or starve certain budgets to then allocate, because that really changes behavior. Because if you continue to allocate that budget, they will continue to use it and go on the old behavior. But, I think human beings are really adaptable. And I think we tend to be good at survival, right? And so, if you force a certain situation, I think people, particularly with the functions of AI today, people will be able to then really try to use it because they don't have the budget for the alternative. I think that's really from a CFO's perspective.

>> Yeah. Well, let me press on that. I find this to be very interesting and maybe even gets to a certain philosophical split, perhaps, among CFOs, because I think, okay, there's one school of thought would say, okay, you keep these very tight constraints, and then you see who can do the most with the constraints, and then you derive valuable information that comes back into your office that then is valuable data. Then, and then there seems to be another school, which is let everyone have really high caps, and we read these stories at American tech companies about the internal quasi token maxing competition. And I think that school of thought would say, look, let the employees go to town, and that is the way that we discover who can actually use these tools productively, etc. But, would you say that that is actually a sort of a philosophical split in CFOs of tech companies in the early years figuring out this optimal path?

>> And and I think that that decision exists within the same company, but it could be two paths that you just described. Because, for example, where we really it's a face of innovation today in terms of where the market's going. And so, we really need to put more dollars in terms of innovation. And you can see that in terms of our spend. So, where we really And that's the emphasis from our chairman and CEO, as well. It's really around making sure that we take advantage of this period.

We're calling this the AI decade within the Novo, and we're in the first year of that journey. And I think this is an opportune time for us because the way we're positioned today, where we have global manufacturing in a fairly fragmented world, where we have a lot of data and production and supply chain concerns, where we are able to be local and be able to supply that to a lot on a local basis and comply to each region's needs, sustainability needs for certain regions, particularly like EU.

Um 100% of our laptops that you like is actually in recyclable box paper. >> Mhm. >> And also, within the laptops, a lot of them are also recyclable materials, as well. And we also have a refurbished effort, given the components constraints today. So, I think from that perspective, we really want to spend dollar on innovation.

Now, as you say, then there's the other pocket. Should the likes of even within finance be able to have a unlimited budget on AI? I think you could probably be sent more sensible in that and say, well, specifically within certain analysis in FP&A or accounting, what was it that you actually used to spend? So, for example, without naming a party, we outsource certain back office function to a service provider. And that's a pretty costly one for us annually. Now, in terms of that capability set, it's really a process flow. So, should we ultimately be able to uh automate everything and be able to achieve savings in that. Now, we are okay to continue to outsource as long as that partner also innovates and be able to share that savings with us as well. So, I think it really depends on where you want to go very specific in each area within your function and your company.

So, in my function, there's so many areas. Planning, tax, FP&A, accounting, treasury, IR, M&A. Right? You name it, corporate finance. So, within all of these areas, where are we going to optimize, be more efficient, productive? And I think that cost dollar sometimes to me is the return is significant. I was just with the head of my tax this morning. And I said, "We If we just pull together and drive the next year's efforts, the tax savings or optimization, right, that we do will more than pay for it, little small trip of 20-something people together."

>> Mhm. >> That's very, very much well spent. But, I think that's just a philosophy. It's less AI-driven, but it's really the philosophy of where that return is. And I think you allocated earlier. You spend based on the ROI generation. And I think today you have to recognize first the AI capabilities that you can have. And I think the other thing we haven't touched upon is really the security within the enterprise of using AI. Right? I think that's a broad topic that we really should touch on. Because today, I think there is really concerns about what we know what we don't know about using external AI within our enterprise. So, I think that's also a concern in terms of how effectively you can use AI.

>> We should definitely talk more about that, but you just mentioned the word unlimited in the context of AI spend. And coming from the US, this is something that we're really seeing and I guess living through at the moment. Like be people talk about the total addressable market being like basically infinite at this point. If you look at the SpaceX IPO, like the total addressable market is now the entire universe, right? I guess my question is, how do you compete against those types of companies who seem to have investors throwing money at them? I know your stock price has gone up quite a lot, so that obviously helps. And then a more broader question, what are the key differences between how Chinese companies are approaching AI versus how US companies are approaching AI?

>> Sounds good. I think there's two very distinct questions there. The first one really with respect to Lenovo's position in the entire tech ecosystem. I was in Davos earlier this year and when I came back I was particularly affected by Mark Carney's speech. And um really talking about the superpowers and the middle powers and of course within the middle powers there are superpowers. I actually used that within our my own management committee in terms of referring our own position within the tech stack.

And I think as big as we are the number one player in the >> power, you mean? >> I I think we're definitely not a superpower. We have to recognize that. But I think from our perspective being the number one device PC manufacturer in the world, having a strong ecosystem and the best broadest portfolio set right across PC,

tablets, and mobile phones, and great partners to Microsoft and Google's of the world in terms of that type of operating system and of course the chip suppliers like Nvidia, Intel, AMD, and Qualcomm. And of course now ARM-based chips as well. And also the Chinese tech tech stack that's also coming up. I think from that perspective, we are there to really enable and I think prior to this stock run, we're really not recognized by the market in terms of our ability to be partners to 2,000 suppliers in the tech ecosystem driving almost \$2 trillion plus of spend and we're right in the middle of it and distributing this across the board evenly in 180 markets. So I think from that perspective we really needed to be recognized. But at the same time we're very humble and I think our chairman really used that word in terms of staying humble. And I think we are because we're also micro and from that perspective most of the profits have been towards the IC companies, the OS companies. If you look at the PC revolution and the tech stack, the OS companies now trillion dollars. The chip companies are now crossing two trillion dollars but have before been hundreds of billions of dollars and we only most recently became around I think 30 to 40 billion dollars but before that sub 20 billion. So I think that's the value chain that you're in but I think we probably did not optimize it as well and I think there's an opportune time for the device makers today to optimize. And of course distribution is is very critical, supply chain is critical, ability to aggregate supply chain to manufacture at scale at cost efficiently and deliver to your customers, to be able to service it, to have provide the security and trust and the after market service support. That must have a value beyond what we're being recognized by the market. I mean I I absolutely believe in that and I think we need to drive that in terms of financial returns from that perspective and AI wave definitely creates an opportunity for the likes of Lenovo and I think we're the best positioned in this because there's no own company like us from the bench to end portfolio and the global manufacturing. So you know, I think that's a very important point. And in terms of how the US and Chinese tech stack today are differentiated or AI companies are differentiated. Today obviously because of the restrictions I know that probably the most efficient computer today is probably from the US chip company. So obviously the chip availability is something that the Chinese air companies probably have to work with. But in that way they probably managed to produce at much lower cost and very efficiently as well. So you see today a lot of chatter about the cost per token generation from the like of a deep sea being and I don't have official numbers. So this is just quoting what I've heard like 1/50 of those in the US. So I think from that perspective, right? As I said, human beings or companies and if you really think about the Darwinism in economic theories, the free markets create the hottest competitors and the most healthy competitors. So I think in China for those that compete in such a There's a word called involution in China, right?

>> We know well. We know well. >> It only happens in China market. So if they can survive within their cutthroat market under the constraints of the chip supply, they are pretty strong. They are pretty strong. So I think that would be one difference I would say. Yeah. >> I want to talk about competition. You mentioned that the deep and broad relationship you have with all these suppliers across the supply chain. And when I think about like when I think about competition within the server space specifically, I feel like there are two dimensions through which you could win. And one of them is of course like okay, you have total supply chain mastery, you can keep inventories low because you have just-in-time delivery of everything and you just have this beautifully efficient supply chain with the world's best parts makers all around the world. And then there's another way you could win, which is your server is just more performant than a competitor's server, etc. And we know other competitors like like a Dell or Supermicro, etc. Those are both strike me as dimensions upon which a company could win the server market. Which to you is more important? Is it the sort of supply chain

excellence or is it the quality of the server itself?

>> I think it's not a either or question. It's fully integrated tech stack. Today, the value to a customer or in a partner, and I think there are less customers today than they are really a partner because their architecture is not driven and the spend is a multi-year spend. >> Yeah. >> So, they're coming to you for the multi-year plan, not just a single purchase. So, we take that attitude and we need to have that attitude with our suppliers as well. We have invested well over 10 years in these factories. We have 30 factories around 12 and each region of the world. We can produce in Europe in Hungary for the European market or the US market. We can produce in Mexico and the US for the US market. We can produce in in Asia for the Asia market. Now, we're building that factory in the middle in Riyadh uh to produce servers also for the Middle East market. And of course, China market, we always have factories there as well for the China market and other markets. So, I think from that perspective, we are really truly global from that perspective. Today, the bottlenecks are not only in getting supply. After you supply, can you then put it together? Do you have the production capability to be able to produce?

>> Yeah. >> Then, do you have the capability to test what is being produced? So, we at Lenovo provide that end-to-end capability to our customer and that ensures them to have the safety and one-stop shop. In fact, our SSG can actually also So, the Solutions and Services Group at Lenovo can also build data centers. So, different from our OEM-like competitors, they can actually build data centers. So, today, the customers can hyperscalers can actually come to Lenovo.

>> Wow. >> With a plan to say I have a plan to have a this site. So people who are not even hyper scalars, people who are land owners, and they have power in specific areas, they come to Lenovo and they said, "Help me build a data center." Because we also have a modular solution that can build this data centers within 9 months. >> Wow. >> In a depending on what you already have available in terms of infrastructure, we can do it as fast as 6 months. So that is very, very fast and that gets the time to market and enable revenues immediately for our partners.

And then in terms of that total infrastructure plus if they have can have GPU compute in the local market, then of course Lenovo can also provide that, particularly with our 11,000 rack liquid cooling capability to be able to service a GPU compute today. Right, and we're expanding on that capability. So today we're the most end-to-end and relevant partner really from that domain. So it's really not from a perspective of really having the supply but and having the technology.

Technology always important, but that end-to-end look, if you think about how people are affecting and enabling AI compute today, they need everything. And of course as Tracy said earlier, I think capital is also important part of that. And you know, it does help that that market is starting to recognize Lenovo in terms [music] of our value add. >> [music] [music] >> Tracy, I'm thinking about the difference between, you know, the server market today and say like the PC market of the '90s. And I think part of the reason we talk about and remember the red button on the laptop is simply because the companies were all unable to differentiate themselves at all product-wise. It was like truly the PC market kind of killed itself because it was just so

unbelievably commoditized.

No one really knew the difference or cared between the difference between the a Dell box and a compact box and a Hewlett-Packard box and a whatever. But it does sound listening to a Winston that like you really can the partner, not the customer. You can really differentiate your offering. >> Yeah, I remember that environment. So I remember like you had to go to Best Buy, right? And you had to basically get a salesman to tell you what the difference was between all these different things.

Yeah. Um okay, since we're on the topic of supply chains and you mentioned data centers as well, give us some color on what it's like to try to get I guess when I think about constraints on the AI buildout, I think about memory, GPUs, CPUs, chips, and transformers now for data centers at least in the US. How difficult is >> Optical connectors. >> Oh, yes, thank you. How difficult is that at the moment getting those key things?

>> Yeah, the the market is very robust because the demand is very robust and then as we're sitting in Hong Kong today but in in the Bloomberg offices, but really there are some really exciting IPOs coming, right? And also capital raises. We're seeing capital raises at such significant levels and that's really all going back to AI infrastructure spend. So really from the perspective of that component shortage or demand-driven challenge on the supply chain, will continue probably for at least two to three years.

>> Mhm. >> But we know that the supply chain is starting to invest, but in terms of you know, some some of the memory specific areas, it does take almost two to three years for a new fab to come online. So I think from that perspective and I think you're seeing bottlenecks across the board, right? And as I said earlier, land, power, not just the transformers, but actual power from the grid. And that's why we're in the Middle East to work with the Saudi government on sustainable, but also cheap solar and much available energy in the local market. So, it's very efficient there, and they have very low-cost energy there that could actually support AI infrastructure from that perspective.

So, we really need to optimize where specifically the power is around the world, rather than just saying, "Hey, everything has to be generated within an area, and that area doesn't even have power." >> Mhm. >> So, I think we really need to be able to do that, pull the global resources together. >> One of the things that we all learned during the pandemic is that, you know, we all learned about how supply chain disruptions work, and the bullwhip effect, or at least, you know, us laypeople learned about these concepts for the first time. Everyone we talked to We recently did an episode with one of the co-founders of Coreweave, one of the neo clouds. You know, everyone we talked to will talk about multi-year-long order books and backlogs, etc. And when I hear that, and I started thinking back, I was like, "Yes, I believe you when you say that, but it could also be a bunch of players double ordering and triple ordering and getting trying to get ahead of themselves just because they're all aware of the shortages, etc." In which case, it does seem like what looks right now like a major backlog and constraints and endless demand could not be as sustainable as people thought. Should we perhaps be worried that some of these endless order books are just in fact companies over ordering because they see everyone else over ordering?

>> I think these are very sophisticated companies in terms of the final off-takers and and uh so, but I think I would point out to the fact that the duplication could happen in the pipeline. >> Yeah. >> But not necessary in a backlog because I think if people have standard definitions of the backlog, by that time it's actually already a sign sign deal and so then yeah. So I think they wouldn't double sign with people.

>> Yeah, okay. >> And then of course revenue recognition is a totally different thing. I think you have already shipped the actual product and you expect to collect on your accounts receivables. >> Okay. >> So I think from that perspective, really looking at the pipeline that is massive in the market, there is likely duplication there and people are waiting to see who has the supply and who can deliver to you, right? But overall, I think the commitment that we see in the market, particularly for data center space, for equipment, for other things, is really a multi-year commitment. And I think people see this as a critical infrastructure. They learn from the internet era in terms of the cloud service providers, how much data can be generated, and this is just going to be so much more. I think it's going to be exponentially more. They're what it is because it will take the existing data plus generate more data and then you need to enable that data, right? And that enablement I think needs additional compute.

>> Mhm. Lenovo has software as well, right? And you have the gaming business. Is that right? >> We have a we're the leading online games Oh, sorry. The game PC company in the world, yeah. >> So, what do you think about the SaaS apocalypse idea? So, the idea that with AI, we're going to be able to replicate whatever application we want, whether it's Microsoft Word or Excel or something like that. Would that eat into your margins or because you own the IP, can you just do it more cheaply?

>> No, I I I think we are not a SaaS company by any means, but I think from the perspective uh of Lenovo or enterprises, I think really as it relates to what we touched upon earlier, your IT department, and I think it actually works for a company like Lenovo because we are not a software-centric company, these tools now lower the bar for us. So, I think our engineers or hiring engineers should be able to create certain apps that should be available or our applications. So, I think from that perspective, there are certain elements of applications that would be able to produce in-house and that would lower the bar for it because these tools are very powerful. But, in terms of the SAS cop popular list, I think it's hard work for me to say. But, I think from that perspective, I believe there is certain data layers that are absolutely essential because enterprises almost cannot get away from from this.

And not nor do they really want to. I think you just need the middle layer to be able to bring the data together and to be able to connect to AI and that's one of the biggest bottlenecks for enterprises today. And whoever can do that, right, would be able to be very helpful. But, I think it's not rather than just purely a software solution, there's probably a consultant element to that as well to help enterprises enable this function. And that But, I think we're we're a long journey from that.

>> You know, I feel sorry for gamers because even prior to AI gobbling up some of the GPUs, right before that it was crypto and I remember they used to complain that the Ethereum miners were buying up all of their Nvidia GPUs, etc. But, I'm curious like you know, this is a one another one of the big meta themes of the time.

Like, in your own gaming business or in the gaming industry in general, obviously the GPUs seems like the highest marginal value you can get from them is with artificial intelligence. Do you see this continuing phenomenon where raw inputs of any sort that might have at one time go to gamers either because just the gamers get priced out or they don't build them for gamers because you can make so much more money selling them into the AI ecosystem.

>> Well, we got a fantastic franchise in terms of as you alluded to the the ThinkPad business >> Yeah. >> is an iconic franchise and and and we love it. Our customers love it. I have actually visited the R&D facility that originated actually from Japan. It's the Yama from the Yama Moto Labs of IBM. So, we kept that intact also with a lot of the people from it from it, but I think the original inventor actually retired some year some years ago, but there's a long history that we're very proud of and we've actually uh maintained that.

In terms of our consumer and also as it relates specifically to the online business online game business which is our Legion business is actually also very famous, right? It's the it's the largest by volume and and people really like Lenovo for the price performance that we can deliver to them and gamers really appreciate that. So, these are amazing franchises that we want to be able to keep and they are a loyal fan base and we have the largest market share in the world. So, we want to be able to protect that. We actually want rather than protect that we want to grow that market share.

So, I think from that perspective we will have to balance between the device side of our business as well as the infrastructure side of the business which I believe they have a lot of demand today. I think it just means that we have to get more share of our supply. >> Mhm. Talk to us about your decision to I guess not tie yourself to any one AI platform cuz I'm sure there are some very very big tech companies out there who would love to have exclusivity with Lenovo and there are also broader arguments out there that eventually we're going to have one AI platform that emerges triumphant in the same way that, you know, we saw Google take over search in the early 2000s. Why did you decide to I guess work with a bunch of different people here?

>> I think today what we see is that we're still in that journey of improvement and you see improvements every few months from the likes of open AI or anthropic or Gemini or whoever it might be. And so I think from that perspective you see a lot of innovation also from the China market. So so from that perspective you really want to be able to be flexible and we think today we're not there to bet on who is going to win. We are there to provide the AI compute. So from now the orchestration there makes a lot of sense for us to allow our customers to be able to go through the Lenovo AI and reach what works best for them rather than them have to figure out and download multiple apps to be able to do that. Right? So I think from that perspective the orchestrator would do that for them and I think that makes it much more efficient both from a memory perspective and compute capability perspective that is on your device. So with that architecture should actually potentially optimize the performance on the device.

>> Should we believe these companies? Let's say in Microsoft. So Microsoft recently announced their flagship model. I forget if it's pronounced Maya or Maya or whatever or MAI. I haven't played with it and they're like and it also runs best in their own custom silicon. How should we read that? Is there in your view a lot of juice to

be squeezed from model silicon alignment or should we read this as at least companies would like to begin having a little bit of a wedge so that they're not so dependent on Nvidia for their hardware needs.

>> I think there's a lot of great things that Nvidia is doing today and I think a lot of people are recognizing. So their their revenues are clearly growing significantly. I think they're probably at the core of driving this current and enabling this current AI wave. >> Of course. >> So I think that continues but of course as cost become to a level where people have to find alternatives, you're also seeing people from a cost perspective making selections, but also probably from as you say the flexibility of having their own architecture. But I think that's really happening with respect to the hyperscalers >> Yeah.

>> who have the scale and also the technological capability and the capital to be able to develop their own >> it about their models really will run better on these chips that they're designing or is it about long-term strategic at least you know, I wouldn't want to say divorce from Nvidia, but maybe you know, wishing they had >> Some time apart. >> wishing they had a pre-nup or something like that. Just like having a little bit of like not so dependent on one company's specific chip capacity and roadmap.

>> I can't speak for for them, but so there's really from external point of view, but but I believe it's probably a little bit of both, right? Because if you were a corporate >> Yeah. >> really deciding on your strategy. I think that's probably one where you especially given that a lot of them are also tech stack focused companies as well. So they probably want to be have more control over their own tech stack. But the other one really is around the cost and also maybe future planning around what they're in enabling. So probably give them more flexibility to be able to do more if they have their own chips in other areas as well.

>> Mhm. So speaking of diversification, if we were recording this episode a year or so ago, we would probably still be talking about AI, but I think we'd be talking about trade and tariffs as well, right? Remember that show? >> I forgot about that story. >> That's right. So Lenovo obviously has a very large and complex supply chain all across the world, but how does I I guess the the general return of more trade restrictions actually impact your business?

>> Yeah, it actually happened very interestingly on April 2nd. So, I took over on the reins. Our fiscal year is April 1st. And I took over on the reins on April 1st. So, the second day, >> [laughter] >> there was the most major event in the history of 100 I think it was 200 plus countries that were enacted on this and we're probably in 180 of those. So, I think it's a major event and given the significance of our business. But, it's it's really about enabling and providing our products at a very reasonable price to our end customers, right? So, I think the end customer suffers if there is inefficiencies added. So, I think that's the most important aspect, which is and I think eventually, right? I think PC products were actually exempt. So, I think there was a conclusion that this is something that was needed, right? For productivity, for entertainment, for a lot of things that people use devices for today. And we're one of the largest providers across the board, whether it's in your pocket every day or on your desktop or at home. So, I think from that perspective, this is probably a very essential thing for consumers today. And I think that's the most important in terms of giving them that product at the lowest cost. You see that behavior really in a lot of probably third world countries or second world countries in terms of

really emphasizing that because it's really about not leaving their citizens behind. It's really about making sure that they have access to digital information and the device is the start of that journey.

>> Right. >> So, that's a very critical aspect of it. So, lowering the tariff barrier is actually essential to ensure that their citizens and any citizens should have the lowest cost access to technology. >> You mentioned something earlier and you're like, okay, on a per token basis, Deep Seek is significantly cheaper than say, you know, the most advanced models from say OpenAI and Anthropic. But, strictly speaking, there is wide agreement that still the flagship models from the American labs are the best models in the entire world and I just don't think there's much dispute about that in either the US or Chinese communities. At least a few months behind. Even perhaps even inside Lenovo.

But do you see in general, and this sort of gets back to the early part of the question about CFO decisions, people not being sophisticated about recognizing that queries don't always have to go to the most expensive, most advanced model. And how much learning is there still yet to do about how to optimize routing of the query such that it goes to the most cost-effective model rather than just slamming the the frontier model.

>> Yeah, I'm I'm going to answer your question in my how I started my morning today because >> Okay. >> I'm in the journey of figuring this out. >> Yeah. >> And I've decided to spend some money with external consultants and so we're in that in journey process. And as I said, you know, we spent an hour and a half with the global team and this is one of the major firms. I'm starting to get a sense from talking to various parties that everyone is new to this game. So everyone, even the advisors who are trying to earn a service fee from you, is really also learning as part of this journey. And that's what consultants do as well, right?

>> ask what you're paying them for? >> [laughter] >> But I I would not uh >> Okay. >> make comments about another industry. >> You love them and they're they're good guys. >> Absolutely. But you know, I think from the from that perspective, that goes to say that if they are the specialists or those that see a lot in the market, then what about specific departments or individuals in the company? So I think from that perspective, everyone is in that journey and it's the people who can be most clear-headed to learn the fastest I think it's probably the most essential.

>> Yeah. >> Right? So I think from that perspective, we need to be able to take out the red tape, enable the process these to work, accelerate. I think timing is and speed is essential in the AI in the in the AI world. And then we need to deploy that where it's sensible. And then as you say, where do you actually allow that spending and the interaction to generate tokens to to really happen. And so we're we're in that journey. But I think we're not alone. I think every company is the same. I'm actually going to a conference next week hosted by a major consulting firm that is a gathering of all the CFOs. But that's in in my regular dialogue with CFOs, I understand I I worry a lot, but I think we're also not alone.

>> Joe, do you think in the future when people like meet each other, they're going to say their names and their nationality and then like declare Yeah, declare Well, I was going to say declare the foundational model that like they use. >> Oh, I was going to say like I wonder if like Ed Bars or like women when they gossip with each other,

>> [laughter] >> they're like, I want a man who is like this much token spend. You know what I mean?

>> I It'll be part of their identities, for sure. >> That's right. Well, okay. So, speaking of rich men, I have to ask at least one CFO-ish question, which is when it comes to capital spend, I'm sure you could justify pretty much anything right now. And there's so much money that's actually going out on the AI build-out. How are you thinking about returning capital to shareholders? Because the stock's gone up a lot, but you know, people can always be wealthier. I'm I'm sure they're into that.

>> Yeah, we paid the in the last fiscal year, which just ended March 31st, we paid the highest dividend ever in the Lenovo history. So, our shareholders are very happy. We're share our success with our shareholders. But we also see a lot of opportunities for growth. And as you have mentioned many times on this the discussion is that there's also a lot of capital needs to fuel for growth. So, we really need to strike that balance between giving our shareholders the immediate cash that they would like, but also the same time they want us to be able to enable the capital appreciation on the stock because the underlying fundamental of the business is growing. So, we need to drive that growth, and if you there's a there's going to definitely going to be a period where you drive that growth.

At the same time we're driving that margin expansion, but there will be quarters here and there that could potentially mismatch in terms of that growth versus the margins. But overall, the long-term trend for Lenovo is a plan to drive growth with accelerated margin expansion. >> Northern Virginia in the US is sort of understood to be the data center capital of America. But you know, it's expanding, there's a lot that's sort of like off-grid in Texas that's happening.

And And then there's also of course anti-data center politics in the US, and we don't know how that will affect the map. What is the Northern Virginia of East Asia right now? Where do you see like which country is it because they have the most whether it's regulations or access to energy. Where is everyone trying to put up data centers? And I'm just curious like anywhere in Asia, is there any sort of equivalent anti-data center politics that's emerging?

>> Well, I think I think Asia tends to be a little bit more flexible from that perspective, but clearly the best infrastructure of scale was in is in China. But I think from that perspective a lot of people do to certain regulations and policies cannot have data centers. So, you have a lot of idle resources today whether it's power and data center therefore accessibility and also supply chain, right? That's not that the world is spending additional money on in higher cost jurisdictions because of regulations that you cannot put it in China. So, the world ends up having some of the inflationary effects is because you're not actually using the most efficient place in the world that you can actually do things.

But having said that, I think other countries are actually catching up very fast as well. So we're seeing even in the supply chain area slightly lower cost. But in terms of the data center aspect and where the power Southeast Asia has been very good and I think what we're seeing in this part of the world clearly in places like Malaysia, Indonesia. I think those are uh uh natural places where people have really uh been looking into. I mean even

even Hong Kong these days I think there's uh uh from the government you would think and also Singapore. So some of these smaller land uh mass areas but I think also have plans and I think right here in town there's also a dedicated uh a few uh players as well and they have to be creative in terms of putting data centers in these higher uh elevated structures.

Uh but there there's there's all all product uh all sort of creativity. And I think Hong Kong has a mixed use in terms of that power generation. But again, that's a small area but I think it's just giving a data set of examples. People are also doing it in Japan. Japan is also another one from regulatory perspective you would do but cost wise and also policy wise in terms of getting necessary permits may not be the fastest. Uh Tracy spent a long time there so so may have some knowledge about this. But I think it really depends on where you can have uh trust and partnership with the local government because it talks it's it involves land. It involves power. It involves imports of specific goods and I think that part really depends on a lot of the local uh efficiencies.

>> Joe, we got to go visit a data center in Hong Kong, right? >> I would love to. Well, I was going to say Hong Kong actually to me makes a lot of sense as a data center location because if there's one thing I know about Hong Kong, it's there's plenty of cheap real estate here and [laughter] there's and there's more than enough space and we know that uh people are very happy with the price of how affordable the residences here etc. So it makes sense that you know I'm being facetious.

>> is actually a lot of land. You just can't build on it, right? >> can't build on it, right. Yeah, so. >> All right. Winston [music] Cheng, thank you so much for coming on Odd Lots. Really appreciate it. >> Thank you, Tracy. Thank you, Joe. >> [music] >> Joe, that was a lot of fun. >> Yeah. >> I feel like we haven't actually spoken about AI from the perspective of a company like Lenovo, right? Like we've never done that before. Yeah.

>> No, we haven't. And that was really nice. And look, part of the I would say two things about AI, which is we're doing a lot of AI episodes. And I would say there's two reasons for that. One is because, I mean, it's just the biggest thing of our lives, probably, in terms of the significance and trying to everyone's trying to wrap their heads around it. So, there's a lot to learn. But it also is in our wheelhouse because it's so physical and because it interacts with supply chains and uh >> Right.

>> And so, you think like, okay, a company like Lenovo, a server company that also happens to uh offer the ability to build complete data centers, not just the servers, who therefore is then interacting with all types of players, is almost like the perfect sweet spot for like an Odd Lots guest, even among all of our perfect guests. Someone who has that visibility on both sides is a really someone interesting to talk to. >> Yeah, I think I said before, there's such a weird tension between these sort of bodyless, faceless AI interface, and then when you think about all the actual physical infrastructure that supports it, like >> It's why I think it's so interesting.

The other thing, too, that I think is interesting about hearing from a Lenovo, and again, hearing from a CFO of Lenovo, is like, look, they're not a hyperscaler. I'm sure they're investing a ton of money in building out their capabilities, etc. But, they're not one of these companies that announces big things like we're going to be spending \$500 billion next year on building out data centers. So, they really do have that dimension where they

have to figure out the Googles and etc. they and the Open AIs they're talking about CapEx and the Lenovos of the world also have to think a lot more about OpEx. And the OpEx component of the conversation, as he said, whether it's with fellow CFOs or the consultant community, it doesn't sound like it feels like it's in day one of figuring it out.

>> Yeah, everything's up for grabs. >> Yeah. >> I also thought the the involution point was was interesting when we were talking about the key differences between China and US AI and the idea that because competition is so cutthroat in China that it just drives your costs down and down and down so that like even if you're not that competitive in China, you would still be competitive compared to Yeah, exactly. >> of funny to think about. Yeah, absolutely. No, I love that chat.

>> All right, shall we leave it there? >> Let's leave it there. >> This has been another episode of the Odd Lots podcast. I'm Tracy Alloway. You can follow me at Tracy Alloway. >> And I'm Joe Weisenthal. You can follow me at The Stalwart. Follow our producers Carmen Rodriguez at Carmen Arman, Dash Bennett at Dash Bot, Kil Brooks at Kil Brooks, and Kevin Lozano at Kevin [music] Lloyd Lozano. And for more Odd Lots content, go to bloomberg.com/oddlots for the daily newsletter and all of our episodes. And you can chat about all of these topics 24/7 in our [music] Discord, discord.gg/oddlots.

>> And if you enjoy Odd Lots, if you like it [music] when you hear from the perspective of Chinese tech companies, then please leave us a positive review [music] on your favorite podcast platform. And remember, if you are a Bloomberg subscriber, you can listen to all of our episodes absolutely ad-free. [music] All you need to do is find the Bloomberg channel on Apple podcast and follow the instructions [music] there. Thanks for listening. >> [music] [music] [singing] [music]