

Nikesh Arora on The Future of Token Costs | Memory Becoming the Moat & Why Enterprise AI Isn't Ready

76 MIN · YOUTUBE · [HTTPS://WWW.YOUTUBE.COM/WATCH?V=V46N1Q7HX1Y&T=4012S](https://www.youtube.com/watch?v=V46N1Q7HX1Y&t=4012S)
<https://www.youtube.com/watch?v=v46N1q7HX1Y&t=4012s>

SUMMARY

Nikesh Arora, CEO of Palo Alto Networks, shares insights on navigating the rapidly evolving tech landscape, particularly in cybersecurity and AI. He emphasizes the importance of adapting workflows to leverage AI effectively, the need for enterprises to rethink their operations, and the balance between product quality and brand value.

- Arora's journey to the U.S. with minimal resources shaped his relentless work ethic and determination.*
- He believes that successful companies must adapt their workflows fundamentally to incorporate AI, rather than just marginally improving existing processes.*
- The conversation highlights the tension between consumer-oriented AI models and enterprise needs, particularly regarding false positives and context.*
- Arora argues that while brand matters, product quality ultimately drives brand value, especially in technology.*
- He discusses the implications of frontier AI models and the importance of understanding their limitations and potential.*
- The need for enterprises to embrace AI-driven decision-making is emphasized, as is the urgency for improved cybersecurity measures in light of AI advancements.*
- Arora reflects on the importance of leadership in shaping company culture and values, which in turn influences employee behavior and performance.*
- He concludes with a focus on maintaining a positive mindset and gratitude, regardless of external challenges or uncertainties.*

I came to the United States with two suitcases, \$200, and I was willing to do anything, anything at all, to make sure that I made a life for myself because there was no way to go back. I was a security guard, I took notes for the disabled, I flipped burgers at Burger King. I had \$200, I had to find a way of paying my tuition.

>> In the hot seat today, we have Nikesh Arora, the CEO of Palo Alto Networks. They have a market cap of 225

billion. Nikesh is one of the most respected operators in technology.

>> I think the long-term token pricing should be 1/10 of what it is today.

>> Mitzvah's, I think it ends up being [music] an accelerant to cybersecurity.

>> In technology, you miss one trick, you can survive. You miss two tricks, you're partly impaired. You miss three tricks, you could be obsolete.

>> Ready to go?

>> [music]

>> Nikesh, last time we did a show, um, I was 22 hours into a 24-hour fast. And I listen back now, and I I just think, my word, you had the audacity to be with one of the OGs of this business and be hangry. Like, I was I was hangry with you and short-tempered.

>> [laughter]

>> Well, the good news is, as I told you earlier in our chatting, I have total memory loss. I have no recollection of what I said and what we talked about, so I have to go back to listen to it. Just happy to be here.

>> It was a great show. We we got on blissfully well. It was wonderful. I don't know what that was about. Um, we

>> I think we should stop doing podcasts remotely.

>> I Oh god, I agree. I I think we should

>> force people to come in here on a regular basis.

>> so so I actually do. And and the nice thing about sizes they do.

But A, the only challenge is they actually sometimes come quite jet-lagged cuz they come like 4:00 the day, and that's a little bit tough.

>> Then maybe you should do it 11:00 p.m.

>> Maybe.

>> Yeah. If you get people from California, it's perfectly nice time for them at 10:00 p.m. or 9:00 p.m. Maybe you need to

>> And maybe I need to adjust. Yeah, I'm the selfish one. Yeah, I was saying Marc Andreessen said that's on me.

Marc Andreessen said I really actually embraced it. Haven't even got to the first question, but [__] it. He said, "You need to embrace how is it all your fault?" If you embrace everything in life with how is it my fault?

Actually, a lot of the world changes.

>> Actually, let's let's spin that a bit.

How do I make it better?

Just change the outlook. It's how do I make it better? What can I do to make this better?

That's how I run my daily and run my life and my company. How can I make it incrementally better today? And how can I make it radically better in 3 years?

>> Have you ever had something that you couldn't make better?

>> Not for the lack of trying.

So, all we can do is try. If it works out, it's great. And if you try a lot, you succeed more often than you think.

>> We were talking downstairs about your personal brand and because making it better.

>> I don't think about it like you do because this is what you do for a living. I think about running my business.

>> Okay, so I think this is fundamentally wrong.

>> [laughter]

>> I'm here to learn. Teach me, Harry. I'm here [laughter] to learn. Teach me.

>> No, but I actually think your personal brand is your business.

>> I think if you build a great product, a great company, people like your product, eventually your brand survives all of it. I think you can have a great brand, shitty product, shitty execution, and your brand goes to hell in a handbasket. It's so I flip around.

>> Do you think that is still the case today though? Like I think brand can be such an exhilarant today, especially in a world where it's just so noisy.

>> See, I spent many years at Google, as you know. I spent actually was chief marketing officer for 5 years. And if you look historically in technology, there are companies which have died who had great brands. Remember Sun Microsystems?

It was a darling, you know, 30 years ago. It doesn't exist. Why? Because their brand went dead or the product went to hell.

What about Yahoo? Remember that company?

It's a great brand. Amazing. Like it's way before Google.

I think it's probably a fraction, probably even two decimal point number versus what Google is. So, I think product helps make brands.

>> But I think But I think those were founded and broke into the public diaspora when there was much less noise. And so, I think if you were to put that today, you need to have a brand first to get into

>> Look, the Let me Let me speak against my own thesis. There's a spectrum, okay? On one end of the spectrum is only have real

differentiated product, in which case the product helps build a brand. Like Google search is now you will Google something, right? On the other hand, it's a commodity. It's water. No one knows why this is called Evian, right? But this is a commodity. Here, only brand matters.

So, it depends where you are on the spectrum. If all you are a commoditized product, yes, brand matters a lot. If you are a differentiated product, then you build a brand on the back of the differentiated product. And you can decide where you want to be in that spectrum.

>> Well, speaking of that brand, I saw your tweet, and I was like, this is masterful. But you tweeted, "Broken up, the frontier model problem is a breadth versus depth problem."

>> Yes.

>> Can you explain that to me?

>> Yeah, look, I I've been really listening to some of your podcasts and I've been paying attention to what happens to market because I want to understand where all this settles down, not just because, you know, I want to understand it, but it also impacts how I build my own business, right? So, you got these phenomenal frontier models, and they keep sort of leapfrogging each other. Every few days there's a net new model that's delivered by OpenAI, or by Google, or by uh our friends at uh Anthropic. And the question becomes, okay, fine, these models are moving in this space, what do I need to build? What do I need to do with these models? And then as we came to that mythos moment, when

everybody was busy chasing mythos, and that was kind of important, you realize even the best model has a high false positive rate.

But for some reason in the consumer space we don't seem to care. You know, I was talking to my sister this morning. She said, "I just went to chat GPT and asked these all these questions. It was very helpful."

So, I guess like what happens is the consumers are way more tolerant of false positives because just it's kind of always the person in the middle, right? There's always somebody who's understanding what the model says and making their own judgment whether they believe the model or not and somehow people get rid of some false positives. Somehow people don't care about the false positives. Some Sometimes people believe the false positive. So, the consumer is highly tolerant on this notion of false positives and doesn't seem to dis- distinguish and it just seems to get better and better.

I literally

had Gemini produce an investment memorandum for something I was looking at. I looked at it. It looked pretty accurate. Like give or take I'd tweak a few things, but it seems passable. So, on the consumer side it's a breadth issue, right? It wrote an investment memorandum for me, which is cool. I would have had to hire a banker and a bunch of investment analysts. It would have taken me days and I did it in 4 minutes. So, the breadth is there, which means it's my go-to place. And as you know in consumer, if you become the go-to brand,

talking about brands, it's hugely beneficial, right? Whether it's YouTube, like that's the only place to go look for streaming video, or Google, that's the only place to go do a search, it becomes hugely sort of multiplicated from a distribution perspective. So, our frontier model friends are chasing the consumer brand and the false positive doesn't matter.

On the enterprise side, false positives matter a lot.

They matter because if you imagine a future where an agent's going to make independent decisions and act on it, you have zero tolerance for false positives.

Now, take Waymo. In my view, Waymo is the biggest agentic product that is out there because guess what? You've replaced a human being called a driver.

Right? All decisions are made by AI machine learning. It decides when to turn, when to stop, what to do.

But think about the amount of edge case training it took to replace that human agent with effectively an AI driven agent. I don't know. Tens of billions of dollars.

So that's what it takes to take one use case and train the hell out of it. And if you think about what happened there, they could have used equivalent of an AI model, but then they built so much context and intelligence and edge case training and proprietary data to make that happen. That data is not available on the internet. You can't stick the next model of, you know, Anthropic into your Mercedes and say, "Okay, drive me home." It's not going to be able to do it. And that's that's the depth issue, right? Because you need the depth of the

context and understanding and the intelligence in around the model to make it useful for the truly agentic use case. So I just think there's this constant tension. The frontier models want the consumer attention because that drives post training for models, it drives the consumer brand of the model. On the other hand, the real enterprise revenue is going to come from use cases that are required a lot more context. The one sort of standout use case we all know is the coding, right? Coding is a universal activity. Everybody does it. So everybody's data is helpful in training the model and that becomes sort of a large enterprise application. Hopefully there's a few more out there. But I think that's the tension that I wrote about.

>> When you think about the workflows and the way the enterprises are engaging with AI then specifically models, what [snorts] extent do you think we will be locked in a frontier model dominant world versus a majority of enterprise workflows can be done with open source and we will be more cost efficient moving towards that most of the time?

>> I think more than half the enterprises are still not getting it right on the use of AI perspective. I think we're still busy trying to incorporate AI into our current business practices. So how do I take what I do today, use a little bit of AI, get marginally more efficient because I don't want to do this the old way. I think the opportunity is to rethink your workflow fundamentally with AI. Um, that's where the true benefit's

going to come. I think the winners in the long term will be people who actually rethink their companies with AI, not people who adapt their current workflows marginally with AI.

>> How can you do that if you're an enterprise? I assume we have we have thousands of CEOs of big enterprises that listen.

That sounds great.

What do I do? Do I do a brainstorming session?

>> No, I think like there there's going to be perhaps two or three different categories. Like one category is let's take the workflows of today. There we have workflows around ERP, we have workflows around sales team management, we have workflows around human resource management. There are existing workflows which have been SaaS-ified as in some software company decided we all have common processes. Let me build a container where these common processes can be marginally customized by enterprise. They can code their workflow into my SaaS application and we're off to the races.

Now, that workflow required a lot of human judgment and human interaction. The software is not intelligent. It's been coded, right?

You define input, you define output. I do the input, I know what the output I'm going to get.

The idea is imagine workflows where you know, in the hiring process most of your workflows are containers. Imagine where AI actually is helping you make judgments. You say put every CV into AI and say these are the 20 people you should interview. This look at the

CV, you should ask this person these following questions. Send a note to Harry saying interview this person, ask the following 10 questions because your three other colleagues only asked the following five. We needed no false positives to human beings. So, AI could be hugely helpful in informing and making the process more intelligent. But that requires us to give up human control and let AI do 80% of the thinking for us. That's not how we're doing it right now. All we're doing is let's take this invoice, let's scan it, abstract the data, put it into AI and say, "Look at that, it's happening 20% faster."

>> Do you think we are willing to give up that human control? You see cases like matters today where I think 1700 people signed a petition that they are unwilling to let it track their mouse and keystroke behavior. We are seeing resistance to giving the data.

>> I would argue they're different points. I think two different points. I think mass collection of data to inform AI is a little dangerous because people see the outcome of what's going to happen, right? And I've heard of companies where people are, you know, using cameras to track people folding laundry and ironing clothes because they want to be able to train physically AI in the future to do those things. So, that part aside, I think on the enterprise side, we can actually run a business much more effectively and efficiently if we decide where we are willing to relinquish control to AI. You perfect example, marketing, right? Anything that is

required to train a marketing model is already out in the public domain. By definition, marketing is public domain, right? If you didn't market it, it's not in the public domain. So, I have the best training data in marketing. I don't need to train an AI model with more marketing content. I may need to train it for tone of voice and what my brand is. And I'm pretty sure an AI model, if I throw my marketing collateral into it, will tell me, "This is not consistent with your brand." If I look at all the things you've been talking about the last 10 years, I think some people have done that. They've actually analyzed earnings scripts of companies and said, "The last 20 years, what has company X done? And when does the CEO start getting away from a topic because perhaps it's not going so well."

So,

>> [snorts]

>> this is really smart. They can do that stuff. So, it's the best marketing training database in the world, the frontier models.

Why do I need 400, 600 people in marketing? Because my biggest problem in marketing is

I have 600 people, but I'm not sure they all fully understand how to consistently deliver my tone of voice, my value proposition, and how not to break my brand by having different collateral in the public domain.

>> You have 600 people in marketing?

>> You have a take on 21,000 people.

>> Uh well, it's

not going to be 600.

>> No, what is it going to be?

>> I don't know. My rule of thumb is that in the next 3 years we'll probably have half the people in G&A type activities in companies.

Things like marketing, things like finance, things like HR, because there's a lot of process management there. And a lot of process management can be made more intelligent using some version of an adapted future AI application, for lack of a better word. So, SaaS applications will give way to AI applications. The difference being SaaS applications have no opinion. AI applications will have opinions.

And that's a fundamental rethink we need from a workflow perspective.

>> Can you just help me out? AI applications will have opinions. What does that mean in terms of how you use them, in terms of the output that they have? What what

>> Everything, right? Like the your scholar, whatever you want to call it, the AI assistant, the AI marketing assistant, the AI HR assistant is going to say, "I looked at your copy, it sucks.

It does not good enough. It's not consistent with tone of voice. Here's what I would recommend." That has an opinion. That will make my average employee much smarter than they were today. Then I don't need so many of them, because they're doing most of the work for you.

>> What would you say to the people that say, "You're wrong. We won't see that halving of those functions, and actually marketing teams will create more copy, more content, be in more places."

>> No, like

I think the places where people could be wrong is how many technical resources we need in the future. I think we need more, not less. I think there's this fallacy people believe we're going to have less people working because AI is going to take over our jobs. I don't believe that. I think what's going to happen is you can't imagine the number of people on my team who want more technical resources, more AI savvy resources, because they wanted to do exactly these things. They're saying, "Well, I've got an amazing project to transform marketing. I've got an amazing project to transform HR." What do you need? "Well, I need more people who understand how to prompt frontier models, build harnesses, bring proprietary data into play, bring modes. I need more compute, more storage, because I want to learn everything. So, I think we're going to need more technical resources. I think we're going to need more sales resources. Because if your product's really good, you need more people to go out there and cover the universe because not enough people know about it. I'm in Europe, I met 20 customers last week. I still see half of them don't know all the stuff we have. I'm like, "Dude, we've been around for 20 years. Why is my team not out there pounding the pavement telling them everything we do?" And the problem is not enough time in a day because I'm too busy dealing with some arcane piece of software at work which I have to go feed. Well, if I had that software be really intelligent, it'd tell me what to do.

>> In terms of wanting more technical

resources,
tokens I would put in the more technical
resources camp.
How do you think about effective token
allocation state? You're seeing very
different camps from your Matters and
Ubers and Microsofts who put in budgets
to your free-for-all
be creative. How do you approach it?
>> Look, I know this whole world of token
maxing has kind of gone topsy-turvy with
the whole conversation that I'm having
with tokens people using.
The challenge right now is
90% of the enterprise employees
are not AI-savvy.
They're not.
They have to learn. I can't send them to
university.
There's no course you can take at any
school anywhere.
They have to be able to learn of their
own. I think we're back to a Darwinian
moment where everybody has to figure out
who's really good. Now, you've seen
people like Brian Armstrong and Jack
Dorsey go out and say, "I'm going to
decimate my organization and I'm going
to start building from scratch." And
they're going to some version of 30, 40%
people because they figured out there's
no redemption. I can't train these
people. I'm going to just find the
people who are going to come in and help
me do this stuff. That's one model. The
other model is sort of gradual. We've
been hiring people only through
hackathons now, right? And we see
natural attrition of 2% give or take a
month and we just replace them with
people who actually are AI-savvy people
who are from hackathons. Give me 12

months, I'll have I'll have sort of transformed 20, 25% of my team. Give me 3 years, I'll have hopefully enough AI savvy people working at Palo Alto. So, there are two different ways to get there. I think part of what you're seeing in the token maxing world is people are learning, people are experimenting. The risk is your smartest employee who knows how to use the AI really well could be using 20 times the tokens that an average employee uses. And if you get into this whack-a-mole moment saying, "Oh my god, I'm going to stop people spending too many tokens." You actually will hurt the best AI savvy people more than you will hurt the average employee.

>> And by the way, I think the best talent will want to go where they will be best equipped with the most expensive frontier models, the biggest budgets. I think that will almost be like an employee benefit.

>> Yes, possibly, but I think part of the challenge is, you know, right now everybody's experimenting on everything. And you have to figure out what is it that I need to build as an enterprise and what can I get off the shelf? Right? If I can get an AI based thinking application that does marketing for me, I don't need to build it. It's a generic problem everyone needs to solve. I can tweak it, I can customize it just the way I did with SaaS applications, but I don't need to build mine from scratch. So, I've made sure that everything my team is building is proprietary to us. This is where do we have unique distinguished knowledge that we bring to bear which nobody else can do on the

outside. Let's put that, let's package it, let's use it. Where it's going to be a generic AI application 12 months, 24 months from now, let's just wait.

>> So, you have a free for all on tokens so to allow your teams to do the best.

>> Uh we have a use judiciously model for tokens.

>> It's not a free for all.

>> Free for all sounds like you can go token max the hell out of it. We have use it judiciously and we keep track of it to see what people are doing. And if we find somebody who's using it well, we won't constrain them. If we find somebody's gone a little off the top over the top, we'll find a way to cap it.

>> Yeah, banning off is said the other day that he spends 300 million on Anthropic a year for his devs and that works out to be about 3.8% of developer salary spend average. Um

if it stays there, the valuations of Anthropic and OpenAI are are grossly overvalued. And if it moves to 20%, they're actually very undervalued. And if it becomes what Brandon at Macquarie said, which is we'll spend as much on tokens as we do on salaries, they're grossly undervalued. Grossly undervalued. I want to talk about where you think percent of developer salary spent on tokens will be in 3 years.

>> Um that's still a narrow lens for me if I very abstract if I step back today.

Um there's not enough compute for what the world is demanding.

Unequivocally not enough compute.

Right?

You can't buy compute.

Compute is costing two to three X or four X more than it used to cost 2 years ago. There's not enough compute.

That scarcity of compute and that excess cost required to build and deliver compute, which allows us to go make AI useful, is causing the constraint and forcing pricing.

Right? Because and interestingly, I think more than half of the compute is going to feed the consumer, which is a fundamentally loss-making entity right now.

Like I don't think any of the frontier models make any money in trying to get you and me to use ChatGPT or you know, Claude or Gemini every day. It's free.

That's a lot of compute. Imagine there's billions of people around the world using for all kind of queries every day. That's sucking away half the compute, which is making no return.

Guess where the pressure goes. The pressure goes on the other half of compute, which is being used for coding in enterprise applications. So, now you're saying enterprise applications coding have to pay until we build transaction models or advertising models on the consumer side because they're not ready.

Now, you could say well, that happened in search, too, right? Google search was around. That happened to YouTube. People used a lot of YouTube, a lot of compute, but it wasn't paying for itself. The problem is the compute requirements and the cost is now 10 X of that when we were in that era versus today.

So, that's forcing token prices to go up. I think the long-term token pricing should be 1/10 of what it is today.

When that happens, you will see that people will consume more. You can decide if 3.8 or 15.8 is not not sure we can tell the answer right now cuz pricing will move very drastically in the next 3 to 5 years.

I think at some point in time

>> Sorry, so in time you think we'll see dramatic reductions in token prices?

>> I think so. I think in the next 3 to 5 years we will see reduction in token pricing. I think at some point in time the consumer use of AI will get constrained by these frontier AI companies because they have enough post-training data, more than they need, and each user is inherently unprofitable in their activities they do in frontier AI models.

>> Do you not think they just build advertising engines like OpenAI is doing now to pay for that business?

>> that's an interesting question. It has to come from somewhere. When I started Google in 2004, we were 2% of the global advertising revenue, and global advertising revenue was estimated between 500 and 600 billion dollars. I think online is about 70% by last count of total advertising revenue. And I don't think the overall number has changed by more than 3% a year or 5% a year.

So, I don't think the total advertising pie is going to increase.

You've already taken away 60-70% of the advertising pie in the online world.

Unless you tell me there's going to

explosion at top where more people going to spend more money in marketing, that money that you're hoping to fund consumer AI from advertising will have to come from current advertising revenues.

So, I don't think that changes the equation drastically to make the consumer profitable. I do think there's an opportunity that AI ends up taking more transaction revenue, which has not been into the purview of AI.

>> Can you explain that to me?

>> Well, think about the marketing chain, right? Like we do advertising.

Advertising is inherently efficient.

What's the best conversion rate do get in online advertising, you think?

>> One, one and a half percent.

>> That's fine. There are still The thing the best of breed is 7 to 10%.

>> Wow.

>> Sometimes, but the average probably one and a half to two percent, which means 85 to 90% of marketing is wasted.

Right? So,

now if you get really smart, you have memory, you have context, and you get smarter in targeting Harry when he's trying to buy something out there in the world, then your conversion rate goes up, right?

If you look at the entire value from the time you decide to buy something till the end, you get the product, you know, you take the case of consumer goods, today, I want to say the cost of consumer goods is probably in the 5 to 8% of total,

you know, list price. The 92% is distribution and marketing.

It's highly inefficient.

So, you could imagine a world where AI makes marketing really efficient, and you get more dollars coming from traditional marketing into the online world because coming in the form of transaction.

>> If tokens get cheaper,

>> Yes.

>> why are we not seeing frontier models get cheaper? We We all thought that this would be cheaper, this would be cheaper.

>> Well, right now they're figuring out that, you know, all your frontier model companies are

value maxing, not token maxing. There is a money at a trillion dollars. At some point in time, they realize, "Oh my god, where's the next hundred billion dollars of compute going to come from?

The financial markets are not going to bear the cost of another hundred billion dollars at a trillion dollars or

trillion and a half because I'm going to need it again the following year. So,

they say, "Well, I got to build a sustainable business model and start showing some degree of gross margin profitability." The only lever they have

is to take the fastest growing thing that they have in their portfolio from an economic perspective and charge us more for it. That's where you get the price of tokens from. I think the price of tokens is high. Now, you can imagine, if the price of tokens is high,

every technology is trying to figure out, "How do I make my compute more efficient, right, in the future?" So,

I'm sure we'll see a whole bunch of advances in the world where memory and computer going to start getting used

more efficiently from modeling perspective. And I still believe I don't need Fable 5 or Mythos 5 to do 90% of what people do with the AI today. Why is the last model not good enough? For certain tasks.

>> I I think I'm wavering here as I think the model over in terms of capabilities and especially consumer and most enterprise demand. Like but the models from 2 years ago were good enough for most of the queries that we asked.

>> right. That's right. The problem is they were inefficient from a compute perspective.

>> Totally.

>> Right? So you're seeing the efficiency come in. The problem is the cost of R&D is now being sort of spent in terms of what the tokens have to pay for. So I think token prices come down.

I think the amount of compute that we need is going to be huge the next 10 years.

I think the frontier AI models are in a position to capture a significant amount of the future economic value of the use of AI.

>> Everyone is fundamentally looking at the stack and yeah this is me being very open as a venture master. That's why the shit's been successful. I'm just saying what every venture master feels. We're all looking at it again.

Oh Jesus, I have no idea where value is accruing.

And you you've got essentially got infra which is kind of top and then you've got models and apps to be totally

>> Yes. Look, infra's making money. Infra's

more expensive than it's ever been.

That's why you're seeing trillion dollar market caps in infra space because of the scarcity of compute and the need for speed.

>> Do you think we're in an infrastructure bubble like people think?

Or suggest?

>> I have a question which I don't know the answer to and you can tell me since you spend more time with people here. I have to go do my day job. Is

at what point in time does physics kick in? And we just can't produce the computer as fast as we want to.

While the infrastructure like infrastructure people are gearing up for large amounts of capacity, large amounts of demand on the infrastructure side and you come to a point that says, you know what? There's only so many data centers we can build. There's only so much

energy we have. There's only so many >> shortages of copper. But you do the Pantalassa, which is building data centers at sea. You got Elon building them in space. I mean, this is

>> So, I think there may be a digestion period at some point in time once we think the demand for computer is there, but the capacity to execute is now limited by physics and the infrastructure players have built up too much capacity for this demand. I don't know when that rationalizes. Maybe it rationalizes and causes us to go think about a different sort of time frame for putting all this computer out. That doesn't take away from the need of compute. We will still want as much computer as we can deliver as fast as we can deliver.

I think some of the model companies have outstripped anybody else's ability to build frontier AI models of that capacity at that speed and the trainings. I think you are seeing perhaps a settling down of who's going to be the frontier model player in the future. The question becomes in the economics, what value accrues to the model, what value accrues to the application layer as you said it. And I think the application layer is probably is a simplistic term because for the first time you have memory and applications. Applications understand context. They understand context as specifically as to what you want, what I want. You're seeing that in the consumer space. That has not yet come to the enterprise space, funnily enough. I think that shows up in the enterprise space, which means the demand for computer and memory goes up on the enterprise side. Like how many I haven't I don't I haven't used all the coding models myself, but over time these coding models have to get really smart about understanding individual context of enterprises and humans. That's how they'll be more effective and more efficient. For that, we're going to still need more computer and more memory. So, I think that'll start defining where the value accrues. I think the value gets shared between frontier models and the context that gets created in enterprise play. I think the frontier models are fully understanding that this is where the gap is. I suspect the frontier AI models, if I had a crystal ball, they will spend a lot more time in the next year or two

uh building memory around consumption.

>> Building memory around consumption. What do you mean, like expanding context windows?

>> More than that. Um

if you look at the consumer interaction that you have with your favorite

favorite favorite frontier model, right?

It's starting to remember, oh, you asked about this yesterday, you asked about

that. Should I take the question you

just asked me in the context of

everything I know about you?

Or should I just limit the answer to as if I don't know anything about you?

Now, that having context of what I said to you over the last 30 days, the last 60 days, 90 days requires you to store a lot of information. Mm. It requires a lot of personalized interaction that needs to have.

Now, if you want to maintain your moat with Harry and Nikesh in the future, the more context you have about me, the easier it becomes for you to give me the answers in the future.

And as you start building context in a user base, you create stickiness, and that becomes your moat.

>> To what extent does Mythos cannibalize a business for you?

To what extent does it make a business for you?

>> Uh

>> [clears throat]

>> Mythos ended up I think ends up being an accelerant to cybersecurity.

>> Yeah.

>> I think what happened when you saw Mythos came out, it demonstrated that all the training we've been giving these models on how

great code is written,
uh the models were able to turn around
saying, well, I also know how to find
bad code.

So, what happens you point the gun the
different way and the model says, oh my
god, look at all this code that you
have. There are so many flaws in it.
As we talked about the challenges,
like every model, it also suffers from
false positives.

So,
if you're an offensive actor and I point
the model against, let's just say, 20 VC
enterprises and I go look at everything
you've done, somebody's left a web
socket open, somebody's done some mess
up in IP addressing, etc. So, it finds
the flaws outside in.

That allows a bad actor to go figure out
how can they daisy chain vulnerabilities
and get into your infrastructure.

It's not good enough from a defensive
perspective because I can't use the
model and say, go take every
vulnerability you found and build a
patch and go patch my system and protect
me. Cuz well, guess what? It's going to
patch 30% things which are not wrong.

Who knows what that's going to do to
blow up your infrastructure. So,

when Mythos came,
we looked at it,

we treated with respect, we ran it
against our code, we discovered it finds
bad stuff much faster than humans can.

We found in 6 weeks what would have
taken us 5 to 6 years. So, we got it. We
ran around, we patch it, but your cloud
code helps build a patch, but you still
have to run it through human evals,
through testing, through production

testing, sandboxing

does this patch break anything in the infrastructure. Only then, after 6 weeks, were we able to go patch everything.

Now,

so what does this mean? This means that every enterprise better fix their stuff faster.

Because if I point the next generation of models against your infrastructure, it's going to find security flaws, security vulnerabilities, misconfigurations, things that you've not been paying attention to. So, it creates a bit of an urgency on the part of the customers improve their cybersecurity posture, which I think generally is a good thing for cybersecurity companies.

>> fundamentally bad. Like when we just like summarize what you just said, it allows you to weaponize bad actors to find holes, but isn't good enough to provide the solutions.

>> Well, look, the solutions are there. Solutions are there. Uh the challenge is getting sometimes getting the attention and focus of the customer saying, "Listen, I got to go fix my stuff because it's important." What this has done is it's lit a fire under the security practitioners around the world saying, "This thing is not good. This is going to weaponize the bad actors. I better make sure my defenses are in place." Now, remember, the way cyber defense is done is it's fundamentally just cyber security is two fundamental things, right? One thing is if it's bad and I'm at the gate, I'll

stop it.

Which means you have to have somebody at the gate. Now, we have 150 million sensors in the world where we stand at the gate protecting our customers.

So, if I can find a way of infusing AI at the gate and taking all these vulnerabilities and finding a way to protect you, I'm good. I don't have to chain the gatekeeper because there's no cloud endpoint agent that exists out there. There's no open AI endpoint agent that exists out there that I can replace Palo Alto or the other people in the space with.

The problem is not at the gate.

What happens is, despite all the perimeter defense you put in, things come in. Things leak in. People make mistakes. People's passwords get breached. You know, there are vulnerabilities people get in through.

Then the question becomes, all right, oh [_] the bad actor is in my infrastructure. How quickly can I find him and get rid of him or her?

That becomes an AI task. That becomes the same conversation we're having so far is I need context. I need intelligence. I need to know what this means. So, creating that context, that intelligence within the enterprise of what this intrusion means and how to protect against it becomes a challenge.

This is the AI cybersecurity challenge. Something we, you know, again, this is not trying to here to pitch my book, but we spent 5 years trying to build that capability inside enterprises. So, in the net net, it ends up being an accelerant. Does that mean I have everything I need?

Not everything. Does that mean I need to get AI models to start helping me? Yes. So, we're going to infuse more AI into our defense infra- infrastructure.

>> Do you think it is good or bad to have government intervention when you have models as powerful as we have them?

>> I think we're going through a discovery process.

Um

I think this notion of guardrails has not been built robustly enough because these models seem to be easy to get past. Remember the early days when you used to read about somebody had this conversation with a model and found a way to, you know, obey the guardrails because it asked questions differently and the model was able to get sort of jailbroken. It became a hobby amongst people and like, you know, they had all kinds of conversation with models. I think it's the same challenge. How do you make sure that you can put enough guardrails around AI model that you built to make sure it's only used for the purpose that you had.

All right, that's the challenge that there is. I I the guardrail needs to get better. To the extent the government feels the guardrails aren't robust enough, it's trying to tell us that it's a national security issue. We need to go fix the guardrails. But I think it's a simple matter of trying to fix and treat the guardrails as a real problem and solve it.

>> Is it possible?

>> I'm hoping it is.

>> I hope it is, too. If you were starting this was I suppose it's one of the

world's best cyber investors. So we might name this cuz he asked some spicy questions to you.

>> Oh, okay. He's got questions for me?

>> Yeah, he asked questions to you. He asked some interesting ones.

But he asked if you were to start Palo Alto Networks for a cyber company again today, starting today, in the age of AI, what would you do differently that you're not doing now?

>> The paranoia I have,

if I look at self-driving,

there are

broadly two or three approaches out there, right?

One was

my car is not going to have a human in it.

It's called Waymo, right? I'm going to keep bounding it, I'm going to keep bounding it, training it until it learns by itself to drive and no human is still behind the

ever when my customers are in it, right?

You see it's out there. There's many cities have Waymos. I saw one down the street from here. So,

that's one way of doing product development. And what you do is every edge case gets discovered, you build training around it, every experience is a learning experience, you build training around it, you keep training it, you keep training it, you keep training it to get to a point of total autonomy. The other version is

I'm going to start taking segments of the driving

and start automating that segment where

I get really comfortable. So, my car

drives 50% of the time, the other 50%

the human gets in on the edge cases.

That's my Tesla, right? Tesla used to drive just the highway for me and now it's getting better at in other streets.

It's slowly getting better, but still

I'm holding the steering wheel very

often. It'll tell me you're not paying

attention. Look at this Look at the

camera, otherwise you can't drive. And I

know FSD gets better, but that's another

way to get there, okay? I'm going to

keep training and fix this. I'm going to

start working the edge cases as my

business continues to evolve.

The third one is I'll infuse some degree

of self-driving into my car because I've

come from the traditional model of

having amazing cars, great V8 engines,

beautiful sleek cars, and I'm not into

the technology aspect. You see them out

there on the street. They have all three

versions out there.

My fear is

do we all need to stop

and start thinking the Waymo way as

enterprises

or is there room for the Tesla approach

to self-driving in our businesses?

Right? Because we have an existing set

of customers to satisfy

who are not going to take kindly from me

saying, you know, guess what? I changed

my product. It's right 80% of the time.

And I'm going to take another 3 years to

train the product to be right 100% of

the time.

Right?

So, my fear is

am I pivoting fast enough in my product

strategy that over time my products

become more self-driving than they are

today?

Or do I need to go faster? And can I get there by automating or AI-enabling certain parts of my product where I can apply the models that are capable to do certain aspects of cybersecurity today and keep doing the others through machine learning and managing S cases edge cases through machine learning or is it time for me to pivot? My view right now is

a you have to have the Tesla approach if you're an enterprise that is building AI-infused capability.

But you can't have the approach of traditional car manufacturers which are trying to stick a little bit of AI and sort of AI-washing their cars and saying, I'll get there eventually.

>> Would you like to do the Brian Armstrong, Jack Dorsey? Like I I often think a good question is what would you like to do

but you're held back from doing?

>> Different courses for different horses, right? I don't think in our business we can go implode the organization because

I don't think the underlying application software is there.

I don't think all the things that we've talked about that we need to get done are ready from an AI software perspective.

>> Uh I don't want to build a lot of software that is proprietary to me for things that should be available for everyone. I don't want to build an AI marketing stack. I don't want to build an AI HR stack. I don't want to build an AI ERP stack. I'm hoping that somebody goes and does that

much more effectively and efficiently for the world at large. Perhaps it's the next iteration of Salesforce, next iteration of SAP, or the next iteration of Workday that is going to help me do that because I do believe it needs to become more intelligent. We talked about it needs to be more AI enabled or AI controlled. Uh I do want to build things that are particular to me where I have all the information, all the intelligence, the context, the memory from an organizational perspective. I think the right way to get there is to hold people accountable. And I run a meeting twice a week now called AI EIO. It's kind of funny. It's like Old MacDonald had a farm. E-I-E-I-O because everybody in my company wants to do AI. So, if you use it as a converging function, as a converging to as a function to do brainstorming across my team. How do we think about this? Why are we building this? I have to AI

>> What happens in that meeting?

>> Everybody [clears throat] comes and shares how they're adapting to the new world of AI. What are they doing from a product development perspective? How are they thinking about it? How are they including agents in their in their in their products? How are they going to go build the back-end infrastructure? How do we think about how are they using tokens? How do we think about the capability of resources? Remember, for me to transform a 21,000 people organization, I have to get the hearts and minds of the leaders to make sure we're all swimming in the same direction or pulling in the same direction. So,

this is my way of ensuring the top 15 or 20 technical leaders in my company are pulling in the same direction. Whatever the direction be. And you know as you know, if there's no expert, then a group of smart people do better than an expert. I don't think there's an expert for future enterprise design yet in terms of here is the blueprint. Like you said, as a VC you're saying, "Holy [__] where's the value going to accrue? Is it going to be in models? Is it going to be in the application layer?" Trust me, we have to have a point of view on that stuff and what's going to emerge in the future before I can start re-transforming my company. Like, how should I build my products?

>> Are your leaders AI built? I interview CROs as well, some of the best CROs, some of the best CPOs.

And in all honesty, the bigger the company, the less AI built, AI massed they are.

To the example of one the other day, public company doing 5 billion plus in revenue, CRO goes, "You know, we don't have that AI talent internally, but we'll we'll we'll bring it in."

>> Yeah, look, I'm

I discovered this in 2004 when I was in Europe around Google Europe. I used to go around meet CEOs. You know, there was this fad where CEOs would hire this 24-year-old sherpa. They were called the the web sherpas or internet sherpas. Right? It's like chief internet officer. Remember those companies had chief internet officers at one point in time?

>> I was eight, so no.

>> You're eight, you don't remember. Well, they used We've seen this movie before.

And sometimes it's fine to have seen this movie before.

And the task of the chief internet officer was to make sure the organization was ready for the internet because I'm too busy in my traditional business and I don't know who Amazon is. I don't know who Google is. So, this wonderful 24-year-old who understands this stuff is going to help be my savior. And then CEOs would wash their hands of all the internet because they have this wonderful team of people who'd be frustrated because they can't get anything done because nobody's giving them attention.

The risk of that happening is true with AI as well. I'm so busy doing what I did yesterday, I have no time to think about tomorrow. So, meet my my chief AI officer was probably a researcher at some amazing university before and has low execution skills. So, until I can get my leadership to understand and agree the extent of the AI challenge and the AI opportunity, we're not going to make progress.

>> What specifically are you not focusing on today because of the burdens of today's problem?

>> Look, everything. When you go talk to a product manager in any large company, I'm pretty sure they have a product roadmap that exists in their heart and in their hands. It's just 6 months or 12 months long. I got to fix all these things. I'm like, what's interesting is in the 6 to 12 month thing, there's nothing called agents in there. How come like the world is talking about identifying everything and your product

roadmap doesn't have that? I'll get to it once I get this done because this is what customers want right now. I'm like, no, that doesn't work. How do I get you to do more agent work? How many people are you going to free up doing development the way you're doing it today so I can use that excess resource to go make new things happen. So, those are all important conversations. I can have them one at a time across 14 or 20 people or I can have them twice a week with them and people demonstrate how they are. And what's fascinating to remember, you have to make sure your leaders are ambitious. You have to make sure they're competitive. You have to make sure they want to win. You have to make sure that they have a learning mindset. When they watch their peers around them do cool [__] they want to show up with cool [__] the next time. So, for me it's getting 14 people together and saying, "Hey Harry, tell me today, what have you done for AI in the last 3 days since I last talked to you in the organization?" And whatever motivates you, whether the fear of Nikesh asking you 3 days again what you did or your sort of inherent learning ambition or it's your team pushing you, you will show up with something. Then you'll see what the other guys are doing. You'll say, "Oh my god, I'm doing a lot or I'm not doing enough." So, it creates a little bit of Darwinian competition amongst them. It creates this urge to go embrace this new technology and I think hopefully I get 14 people fully motivated and then they go to that to the next set of people cuz

I need to transform from the top down,
not from the bottom up on this topic.
There's a bottom up experimentation of
course, right? People using tokens to
see who's really good at that allows me
to find the best talent. So, you got to
find a way of transforming 20,000 people
over the next 2 years in that direction.

>> Bottoms up, top down, you got to get
into these organizations. That that
sounds normal.

>> [laughter]

>> I was just you know, I wish I was a VC I
could have said that in the podcast and
talk about stuff and

>> Clearly I need to get out more when
that's

>> [laughter]

>> Um I just sit in this dark room all day
in a cave. I'm sorry.

Okay.

The question is, how do you get in
effectively and how do you get
implementation and adoption done well?

I've had guests on the show say before,
you cannot do enterprise adoption
without FTEs today.

And then I've had my time come on the
show,

our friend from Factory, and say,

"If you need FTEs, you have a [__]
product."

Bold. And I love my time. And I

I think it was a great clip. So,
grateful for the virality that that came
with that. Is that that one viral? That
one did very well. Yeah, Shyam from
Palantir then obviously chimed in. And
my job is to create a discussion. Um,
what is true and what is right? Do you
have to have FTEs to sell into
enterprise?

>> So, I think what
what is true is
we've only been chasing the enterprise
dream for AI for the last, what, 12
months?
At best.
If you think about
everything that happens on a weekly
basis, we see new things come which we
don't quite fully understand and grasp,
right? We're all busy trying to get our
arms around LLMs and how they're going
to be great for chatbots to talk to our
customers in enterprise, and suddenly
agents showed up. So, it's like, "Oh my
god, I got to figure out agents. They're
going to start working internally.
Agents are going to do a lot of stuff."
I'm pretty sure you can still have a
agent fest and have everybody tell you
what an agent is. You still have to walk
out and say, "I'm not quite sure that
his agent or her agent is the same as
what the last guy said." So, because AI
is moving so fast, I don't think the
products are fully there yet. Like, the
enterprise products at the application
layer don't exist in their entirety
because we haven't been tested against
the enterprise ask. So, FTE is a short
form for saying, "My product's not fully
there because it's evolving as the
technology evolves.
I'm going to send some people across who
are going to sit in your office and
build my product
while I adapt to your needs." That's
That's is, right? That's what we saw
from Palantir. That's what we're seeing
from all these companies. So, what
you're saying is, I'm going to send my
product engineers or developers to your

or enterprise, and they're going to build my product. You get it if you do it right. Again, FDEs are different versions. Some people are just trying to get you to consume AI, which is actually not an FDE, it's just a technical sales consultant who's trying to help adoption. On the other hand, an FDE truly is somebody who actually brings the code back from a customer side and goes back to your product and say, "Listen, I built this at the customer side because they had this need. We should incorporate this into our product because everybody's going to need it."

That's an FDE in my mind.

So, I think FDEs are needed for the short term because, remember, all the enterprise AI startups are hungry for revenue. For some reason, we've created this notion that don't worry, just keep keep selling it. There's a huge sort of pent-up demand around AI applications. Sell it before the product is fully ready. That's what you're saying.

>> Do you think that's right?

>> I think that's the case. I think as we think things evolve in the next 12 to 24 months, people will switch from one set of products to another because something will emerge as a better product. Look at the coding conversations, right?

How many coding companies have you heard of the last 24 months? I think we had Wind Surf, we had Devin, we had which is now Cognition.

Wind Surf got sold. So, those are the early guys in coding. They don't exist in their then form.

Now, you've got Codex and Claude and

Antigravity.

You got Factory doing SDLC, you've got Cognition doing SDLC. So, you can see as the market evolves, people who have concentrated on different parts of the value chain around coding are getting formed. The product is getting more and more formed over time. And who knows in 2 or 3 years who's going to be the leader in that space because the product is not fully ready when you start.

>> Do you want to make a bet on who will?

>> No, you do that.

>> [laughter]

>> I just need a good one that my teams can use. I don't need to make a bet.

>> Uh I I love that and and I agree. Can I ask you, how do you choose who you decide to help?

You You know, I spoke to your daughter Aisha before, and she said one thing that not I should lots of things that many people don't know about you.

Um

but she said one is you help a lot of people, a lot of founders, and you ping them.

>> Yes.

>> How do you choose who you ping and who you help? Mattern obviously being one of them.

>> Well,

my current paralysis, as I told you, is this market is moving so fast that based on what you read, that Open Cloak comes out, suddenly there's this thing that people going to have agents. There was a moment, if you heard If you saw, there was agentic browsers, remember? You don't hear about them much, but there was a moment when everybody was going to have an agentic

browser, your browser was going to be your computer, and that's going to be sort of do all the agentic tasks.

When I hear about these things, I'm trying to assess which one's going to work.

If it works, how does it impact my product portfolio? What do I need to build in anticipation of this technology becoming mainstream? And that window from the idea to execution is shortening in the air world, as you can see, right?

Like, the way these companies are coming out and getting formed in 12 months and 24 months and getting a 100 million ARR is probably the fastest ever, which means if that's what my enterprise customers are using, I have to figure out how to secure that stuff. Well, my team doesn't fully understand all this stuff. So, I'm really listening to podcasts, listening to people, watching people tweet, watching people on LinkedIn saying, "This is an interesting technology." Helping the founder. So, my first step is to ping somebody who's doing something interesting, which I don't fully comprehend. They seem to be getting to a degree of success, which gets me gives me a feeling this could be something relevant. Perhaps not this company, but the construct that they're working on, the concept they're working on.

>> I'm an ambassador in a world of embassy, in a world of uncertainty, you go later where there's more certainty. I we We about this downstairs. Um you know, I have that luxury in terms of a flexible mandate to do that. You have that luxury, too, in terms of the benefits of scale and acquisition

budget.

Can you not just sit on the sidelines and wait for the right things to percolate and then buy them at a billion?

>> Yes and no.

Um,

yes, we can wait.

Uh, that doesn't mean I don't need to learn.

If I'm not paying attention to eight different players in the space, which I'm sure you do, too. I'm not paying attention to eight of them, trying to see who succeeded why, what did they do wrong, what are the guys who got it right too. It's very hard for me to assess what made it work. Was it a fundamental, like, it was a bad idea?

The technology is bad. Like, agents are not good. This agent is not going to work. It could be that, or this company didn't implement right. The agents are still a phenomena. Somebody else is going to execute right. I need to understand the underlying technology for sure for for to make sure that my team is thinking about it and do we adopt it adopted and secure it in the future, right? We bought a agentic AI company gateway to um

6 months ago. Didn't cost a lot of money.

But, I figured out saying, "Listen, if everybody is going to agentify the enterprise, how are we going to know how many agents you have running around the enterprise, how do I keep track of them, how are we going to govern them, how are we going to secure against them?" So, I said, "The only way to do that logically is to find a way to aggregate agent

traffic somewhere.

If it goes through a certain gateway, a firewall, or some router,

I can watch all the traffic and I can stop an agent from acting. That's the only way it works.

So, I said, "The first thing you need to be able to do agentic security is to have some sort of a gateway." So, we bought a gateway product. Now, I got it at the right price. If I wait and look at what's happening now, suddenly people are waking up to the idea we need some sort of a router or gateway that all traffic needs to go through because of optimization reasons, because of routing reasons, because of token maxing reasons, which you have paid double.

Maybe.

It wasn't a big price, but, you know, I could have paid double. That's not the point. The point is Look, things I buy, either they're going to help me 10x or 100x, or they're going to fail spectacularly. It doesn't matter if I paid one or two x at that point in time. Of course, I shouldn't be paying two x and having it fail spectacular all the time, but the one two x doesn't make a difference. The one is to 10, one is to 100 is what you do. That's what we'd like to do as well. Not from an economic return perspective, from a business value perspective in our in our business.

>> Are you more involved in Corp Dev today than you've ever been?

>> No, I've been always more involved in Corp Dev. This is not a problem.

>> normal?

>> I think I'd say I'm more involved in trying to learn what's happening out

there from a technology perspective than I've ever been, because the stuff is moving so fast. And if I don't have a point of view, and if I don't encourage my teams to pay attention to it and we talk about it, I think there's a risk we miss a trick.

And if you miss a trick, you know, in life, in technology, you miss one trick, you can survive. You miss two tricks, you're probably impelled. You miss three tricks, you could be obsolete.

>> A lot of SaaS providers are feeling obsolete today. Their share price is telling them they're obsolete.

Do you think that the majority of SaaS vendors have been oversold, or do you think that is an accurate reflection of where markets are moving?

>> I think what the market is telling us is that

the system of work or the systems of record will see a reimagination of workflows, as you and I talked. So, going from software that doesn't have an opinion to software that has an opinion and expresses an opinion and also does a lot of work for the human, so the human doesn't have to do repetitive tasks. I don't think those AI applications have been created. I think the SaaS versions exist. We all use them. I think at some point in time we'll see AI applications that do a lot of the task and workflows get reimaged.

I do think a lot of SaaS has built a lot of analytical capabilities to sit on top of the systems of work and systems of record. I think it's a lot easier to abstract that data into some large data lake and have LLMs analyze that data for you and give you the answers. I think

the analytic world is getting reshaped already. Where you can see people like Snowflake or Glean or Databricks, all these people boast enterprise data lakes where you can bring the data and run LLMs against it and get you much more synthesized analytics and outcomes than you ever had before. I think the third question which we started off with is like people are not sure how many people are going to work in these enterprises in the future.

So, if you take the the confusion of how many seats are going to survive in a SaaS world, you take the confusion around analytics are going to get done differently, and system of work gets reimaged. So, I don't know what the right valuation is.

>> done differently? So, I'm again, I'm disclaimer, I'm a podcaster.

>> Oh, you get it. Yeah, yeah, you're successful investor managing lots of people's money. That's true.

>> Uh but disclaimer, podcaster.

>> Got it.

>> Uh I understand the seats question. I understand the workflow. Can you help me understand this?

>> Well, if you look at most SaaS software, right? In the past many years, once you're fully deployed at a company, the company says, "Listen, I've got all this cool data about all your employees in my HR system and I can help you get more insight in HR system." Or if you take a Salesforce, they have a Salesforce marketplace with 300 apps you can use which are analytical apps that feed off your own system of record go-to-market data and helps you analyze that data.

>> Do you
Do you Do you know Neill Mehta?
>> Yes, of course.
>> Yeah, I love Neill. I think he's one of
the most phenomenal people. Um
I'll never forget him being in London at
Christmas time and it was like I think
it was Christmas Eve or the day before
and he was spending time with me going
through my pre-seed portfolio.
>> Right.
>> And I think that just shows the hunger
that he has for learning the day before
Christmas.
>> He was a catch the cool early company so
he can make lots of money like it.
>> But just so intent on finding the next
thing. Um and he always says the one
question is like, are companies' best
days ahead or behind it?" And that's a
very helpful one.
>> good question. Good question, yes.
>> Our sales force's best days ahead or
behind it?
>> I don't know. That depends on how they
execute from here on.
>> If I were to paint a bear case for you,
>> Mhm.
>> what would that be?
>> The bear case is we don't get this
transition right of the world going to a
AI-first
future.
Because look, these false positives will
keep reducing over time. Agents will
become a real thing.
Agents will do a lot of work for humans
which humans have been doing manually in
the past.
All that needs to get embodied in your
product.
If I can't make that transition happen

with my team in the next 3 years, yes, there's a bear case because somebody else will build a better mousetrap.

>> And the bull case is you understand it better than any other security provider, and you become the default security >> bull case is that we get that transition right.

There's always already a trend in our favor underlying that where people are realizing they can't have 40 to 60 cybersecurity companies that they have to manage themselves. So, we've been driving this trend of platformization already for the last 24 months, 36 months. We already see the fruits of that where people are saying, "You know what? I don't want 40 people solving my problem. Let me put Palo Alto and solve the problem that 20 different companies do together on one platform." Now, the good news is because we are all coming to our senses and saying, "Yes, we need a lot more enterprise context, enterprise data. It has to be stitched. It has to be seamless." That's what we learn with our current proposition. I just need to bolt on the right AI not bolt on or embrace the right AI capabilities in that stuff.

>> Does the platformization remove the ability for venture scale returns?

>> Does the platformization

>> Remember, I need like \$10 billion companies. Like, this is the big thing that I think most founders still don't kind of fully comprehend. It sounds awful. A billion dollars doesn't do it anymore. It needs to be 10. It needs to be 20, 30. With the platformization, I can get a billion-dollar exit to you, hopefully. Please buy any of my

companies for a billion in cash. I'll give you the catalog. You can take them.

Uh

uh

>> [gasps]

>> but I you know what? I'd pay 10.

>> I'm not going to fight against innovation. I think there'll be venture scale returns in cybersecurity because remember, we're the most innovative industry in the world.

The bad guys are always looking for a new way in. They're not saying, "Oh, I exploited that 2 years ago. Let's try it again. Maybe somebody hasn't deployed a patch against that. Sure, we fixed that one." You got to go find a new way to attack people.

>> to come up with an innovative way to hack into it. Let's just try to do Westlife again.

>> Exactly, right. So, it's highly innovative. There are new attack vectors. People are going out there trying to chase them. I'm not going to build everything myself. People will build great stuff. And sometimes people will great stuff and build a platform around it. And that's fine.

Remember,

we come to a different vantage point.

When I started Palo Alto, we were less than 2% market share in the entire revenue of cybersecurity. We're closing in on 8 or 9% right now, right? There's still a lot of room between 8 and 9% to 20 or 30 or 40.

That means there's still 60% market cap out there to go enjoy in different companies. And that's not all going to be existing players, including us. There is room to build companies which have

tens of billions of dollars of market cap in the next 10 to 25 years.

>> [_] enormous market, eh?

>> It's beautiful.

>> Yeah. Wow.

>> Well, think about it. The S&P, what percent of the S&P is tech now? Compare that from 20 years ago.

>> Uh year-to-date gains like 86%

>> Forget the gains. What's the total of the total S&P market cap? What percent is tech? And what was that 20 years ago? What will that be 20 years from now?

It was less, it'll be more.

>> [laughter]

>> Yeah, you know, I I 100% I'm aligned in that area.

>> So, so the entire tech tech space, like, you know, what do you call marketing tech in the future? Or marketing spend or tech spend? What do you call HR tech in the future? HR spend or what do you call the spend on all the tokens which replace repetitive human tasks? All becomes tech spend.

>> You said a word, bad guys. Is China in many people's eyes. And we see a huge amount of incredible open source models, which are being used extensively state a much cheaper cost. Do you think the proliferation of Chinese open source models is something to be concerned by or is an inevitable feature of a burgeoning ecosystem?

>> So for a second, let's play the thought experiment.

Take the word China out for a second.

Answer the question.

>> Do I think open source models

>> Yeah.

>> No.

>> You think open source I don't know what

they say.

>> I don't think open source models are uh dangerous. Uh

>> Right. So does it matter where they come from?

>> Yes.

>> Okay. So you're not worried about open source models, you're worried about Chinese open source models.

>> 100% I'm not saying I am not.

>> Remember there was a large tech company which also had open source models for a while.

>> Sure.

>> Right. So

it's interesting to watch open source models. I think the question becomes in the future

do we end up with you know, horses for courses? Do we end up with models that are very task specific and they're helpful in certain tasks and do we then always need to use this mega frontier AI model for everything? And you already see that with 11 labs and you know, the voice models that are out there, which are specific to a task and they probably do that task better than what the frontier AI model does. So we're talking if you believe the world bifurcates into many task specific models, which are going to be useful for that task. That task specific model could be better training across the depth in a vertical space. That's going to happen physically AI for example, I don't think physically AI will be as easy as having a generic frontier model because there's no consumer use case of physical AI.

Right? It's a depth use case only.

The question is is your physical AI model that helps you fly planes going to

be the same physical AI model that helps you drive cars? Most likely not. Will it be the same physical AI model that does robotic manufacturing? Probably not. So you're going to see depth in these models. You're going to going to see a world of bifurcated models with some sort of orchestration layer as we've talked about and you're busy finding orchestration layer companies that can allow you to pick the best model for the right task. That those orchestration layers have to get smarter and smarter. They have to understand the context, they have to understand the memory, they have to So, the question is, do I store my memory and context in the orchestration layer or do I so store that in the frontier model?

>> Which one do you think it will be?

I know I'm the investor, I should know, but I don't.

>> No, I I think that the challenge is right now, the frontier models know this problem and they're aggressively moving to incorporate memory and context into their models because they understand that's the mode.

And

the challenge is you have to pay for twice. If you say, "No, I don't want to use your memory and context." the model may not be usable. If you use an orchestration layer, the orchestration layer today is not as well funded as these models.

The risk is you end up in an architecture where the model has a lot of context and you cannot be model agnostic, you actually be model captive to get maximum efficacy and value for what you're going to get done.

Right? It's not like
it's like you have a choice.
You have to go all in on a model or you
can't go all in on a model. You can't do
with one what you can do with the other.
If you want to do with the other, you
have to redesign your entire application
that is deeply embedded with the
capabilities of the second one.
So, in the world of bifurcation and
horses or
sort of courses,
I think open source is a good thing
because it allows you to play the cost
curve. I don't need the smartest model
to do the smartest thing.
So, open source is good. Whether it
comes from a certain country or not, the
question becomes
you know, what
what backdoors are you worried about
that these open sources model models
have? What are you worried about that?
And that's true for any nation-state,
right? If there's a nation-state
sponsored open source model, what are
the backdoors? Can I get in? Does the
model wake up one morning and it's got a
sleeper agent in it and starts sending
all the data somewhere else? Those are
questions. Those can be secured.
That's why you come to Palantir to help
you secure the models.
>> Who is going to secure your backdoors?
>> [laughter]
>> I don't know. What what kind of ID you
have?
So this is not a hangry night. This is a
different kind of night that I'm dealing
with.
>> This is Monday morning.
>> Yeah,

>> [laughter]
>> all right.
>> It's just terrible.
Um
>> To to
>> What is the best time for me to show up
here?
>> [laughter]
>> I am in your time zone. I'm well rested.
I'm not jet-lagged.
>> [laughter]
>> The quote from the the catch only secure
your back doors,
not the
>> mind is going
in the wrong direction. The only
three things you seem to have
caught on to
>> I know. I've got I've got two questions
and we'll do a quick fire. One is um
with the incredible success you've had,
um
you've made a lot of money. Um
and and the just the question I have is
like, what does no one know about having
money that they should know? Like one
thing for me is I've become much more
impatient. We have very different You're
You're far more successful than me. I've
become way more impatient. No one told
me I'd become impatient. I I'm used to
good quality of of everything. And now
when it's not that, I'm very pissed.
I don't like that in myself. I But no
one told me it would happen.
>> fix it?
>> Therapy.
>> [laughter]
>> This is the common Western solution.
[clears throat]
Yes.
>> Have somebody else tell you

>> [laughter]
>> You've gone back saying, "Oh,
>> I must feel better now because I told
somebody that I was going to therapy.
>> told me I should put myself first more
often.
>> came back to being impatient because
that's how you put yourself first. More
important. Believe in yourself, right?
Is that what you're supposed to do?
Optimize for yourself.
>> myself first and then it was and it was
your dad's fault.
>> Oh my god. Is that what your therapist
told you? That must be British. Okay.
>> Um but no one told me that. I would I
would kind of wish they had done. It
would What does no one tell you about
having money that they should do?
>> I think it's not about money as much as
it about it's about success. Remember,
we all follow Maslow's hierarchy.
I came to the United States with two
suitcases, \$200, and I was willing to do
anything anything at all within reason
as the right side of the law to make
sure that I made a life for myself
because there was no
way to go back. I was going to use a
different word, but I'm not going to use
it because you'll go crazy again. So,
>> [laughter]
>> there's no going back, right? It was a
one-way ticket, which I did not have
I had no recourse. So, I was willing to
do whatever it took. I took notes, I,
you know,
became a security guard, I tried to pump
gas for a weekend.
>> You became a security guard?
>> Yeah, when I came to the United States,
I was a security guard, I took notes for

the disabled, I flipped burgers at Burger King. I had \$200. I had to find a way of paying my tuition.

>> Which was

was one quite transformative to your mindset? Did you hate them? Did you love them? Like, what's that like?

>> It had to be done. It's karma. This is when you come from Eastern philosophy, it's karma, right? It's destiny. This is what you need to do to break your destiny.

So, you do that. So, you don't worry about what you have to do. Now, at that point in time, there was no

>> know you were going to be successful?

>> I don't know. Who knows? Nobody knows they're going to be successful. You just come in and do your best and hope for the best and see what happens.

So, that's very Eastern Eastern philosophy, right? You be You believe in karma, destiny. How do you manage billions of people in the world? You make sure they believe in destiny. If they believe in destiny, they'll say, "Oh, this must be what was my destiny in the end. I tried my best. This is where I ended up." That

It's better than your therapist. Just like keeps you centered. Saying, "Okay, I tried my best. I gave it everything I had, but perhaps this is what God intended for me."

You find that hard to believe.

>> If you were my therapist, I don't think I could afford you, though. That's the point.

>> This is free. You're getting to getting free.

I'm not sure I fully embrace all of that, but that was where I started. So,

if you start perspective, over time, you climb up Maslow's hierarchy. It was about food and shelter, and then it became about ambition, and it becomes self-actualization.

Right?

Conceptually, in Maslow's hierarchy, and you get to a certain amount of money, then you decide there are some things I don't have to do anymore. I don't have to be a security guard. I don't have to flip burgers, right? But that very quickly goes up for the saying, "I don't have to tolerate certain things that I tolerated in my life because I don't need it in my life because I don't have to adapt to the circumstance because I can walk away.

>> Do you ever get worried that the willingness to walk away makes you softer?

>> The willingness to walk away makes you softer. No, actually it's the other way around.

The willingness to walk away make sure you optimize the outcome. When you negotiate, if you're fully vested in the outcome, you fold at some point in time saying, "Well, I can't let Harry Stebbings walk away cuz Harry walks away, I have no deal." But if I say, "You know what, Harry? It's going to be these terms or no terms. I'm willing to walk away."

Then it depends, a battle of wits, right? So it's just Harry want it more or do I want it more? Does it make you softer?

>> I don't want to make it political, but isn't that Donald Trump's like The Art of the Deal, like leverage?

>> I don't know. I've not read the book,

so.

>> It's quite a good book, actually.

>> Is it? Good. Like at the end of the day, I I don't think being willing to walk away makes you softer. I think being willing to walk away makes sure that you understand the pros and cons of what you're dealing with. It makes you understand whether you should spend your time or there or not. It just makes you understand whether you can get an outcome that is useful for you as well as the other person. You know, we have a lot of choices in life.

Once you have the amount of wealth you have.

>> I'm on the forecast, sir.

>> Easy, easy, easy, easy. You You said that a few times.

>> Final one and then we'll do a quick fire. I care a lot about kids, actually.

I love kids

um and I want to be a really good father when I am one.

You're a CEO of a public company that is incredible. You've had an insane career.

And I've had the pleasure of meeting, you know, one of your children. She's amazing.

>> Cheers.

>> She's amazing.

What advice do you have for me on lessons on how to be a great dad, but also not losing an inch on work. I'm not willing to sacrifice much on the work side.

>> Yeah, this is the hardest problem in the world.

I think there's what, 20, 30 billion people who have been born since the beginning of civilization.

Yet there is no AI that can train us on

what we need to specifically do to
create the outcome we'd like to create.
It's way too many variables, right?
There are all kinds of people in the
world and I'm sure their parents Some of
the parents are amazing, some of the
parents are not as amazing. So, I think
part of it is
you can do your best from your
perspective and I think kids absorb a
lot
by watching you, your work ethic. They
watch your values, they see how you
interact with them because
my daughter probably has a better sense
of who I am as a person
than anything I can tell her because she
spends time around me. She sees me
interacting in every sort of micro
situation, what makes me impatient, what
makes me patient, what makes me do
certain things. At the end of the day,
your child believes that you have the
best intention for them, I think that
goes a long way.

>> I totally get you. It's I had a guest on
the show and they said, "Watch
National Geographic if you want to be a
good parent." I said, "What?"
He said, "Look at the elephants.
The children follow. And so, if you want
your child to be nice to waiters, be
nice to waiters. If you want them to
work hard, work hard."

>> Yes.

Well, that's true in organizations, too,
by the way. Organizations take on the
form of the leader. I'm pretty sure if
you close your eyes and you rattled off,
you know, five or six attributes of
a company
and said, "This company has a founder."

And say, "How do you compare the company's cultural values vis-a-vis the founder?" And you'd find a remarkable resonance between the two things.

Companies act because remember, the organization is trying to please the founder because they figured out that's the way to achieve success. If my CEO is impatient, if my CEO is exacting, if my CEO is ambitious, my CEO suffers no fools, gets stuff done, then that must be what they want to reward. So, you suddenly find if you This again depends if he has the right values or not. And you told me a story about a guy who had different values and they had to shut the company down. But, if he has the right values, people will watch your behavior and want to emulate your behavior.

>> [snorts]

>> I want to do a quick fire because otherwise I'll take up all of your time.

Um

what is a belief that's held by most top investors and founders in Silicon Valley today that you think it's wrong?

>> My concern would be at this point in time

given the um base at which technology is evolving

given the

uncertainty in terms of what's going to work, what's not going to work

and where there might be too much euphoria and a bit of FOMO going around in terms of my god, if I don't invest in something that's interesting and the right founder, I'll be left out. And because people have seen this happen.

Look at what's happening in Anthropic, right? If you missed the first round,

the second round, the third round, the fourth round, the fifth round, then you look like a guy who's got money. Now, you had 20 years to invest in SpaceX. You had three to invest in Anthropic. That piece is fundamentally different. I'm sure as many people are happy that SpaceX finally went public, as many people are probably sitting there saying moping saying, "Damn, I should have done the Anthropic round 2 years ago when they showed up on my doorsteps." I think there's a lot of FOMO coupled with euphoria on the other side and I think the risk is that we think every company that's going to show up now is going to be the next Anthropic, so we better get into it.

>> My next one to you is any moment, any board meeting, what was the biggest oh [__] in a board meeting?

>> I got a very interesting an insight from one of my board members. Um you know we're prolific buyers of companies because I'm constantly paranoid that we haven't built it, somebody else is going to build it, so we better go acquire it and find the team to go get it done. And uh it was one particular acquisition it took a lot of effort to get the founders to table, get them to agree, grind through due diligence, figure out whether it's going to work or not. And it's it's a substantive amount of money, relatively speaking, hundreds of millions of dollars, close to almost a billion dollars. And I call one of my board members and I said, "Hey, what do

you think about this?"

He said, "You're calling me. You don't call me all the time about all the acquisitions you do. So this must this one must be different." I said, "No, it's not different. I'm just thinking hard about it. It's taking a lot of effort." He says, "Go for a long walk. Ignore all the effort you put in."

He said, "Because sometimes what happens is

you confuse effort with wanting to get the outcome. Cuz I spent a lot of time and effort trying to get it, then you feel like when you get it, you better take it because you put all the effort in." And he says,

"You haven't spent a dollar yet. You just put in 3 months of effort. But remember, once you put the dollar, then it becomes yours. It's your job to make it successful. So, you still have one more chance to decide if you want it or not." I go for a very long walk and say, "If this walked in the door right now, and there was zero effort involved, all I had to do was write the check, would I take it or not?"

>> Forget the sunk cost.

>> Yes.

>> Have the same in the investing business. You spend so long

>> Yes. Yes. You spend a lot of time. You're like, "Oh my god, I'm the one getting the term sheet, and nobody else has it. I nailed it. I've beaten out eight VCs to it." The question is not how many VCs you beat to get the deal. The question is not how many VCs you beat to get the deal.

The question is, "If this deal Can this deal stand on its own merits, and would you invest in it if there was no competition?"

>> What's the best advice you've ever been given?

>> The best advice that the really old man gave me on a flight once was, "You know, life is simple. If you wake up in the morning, you're really excited about going to do what you do for a living, you're blessed. And if you're done after a long day and really excited to go home to your family, you're blessed."

>> [snorts and laughter]

>> I do want to end on a note of

>> like that?

>> I do. I do. And I actually tweeted last night, "I hated school when I was a kid. Sunday nights was the worst."

>> Yes.

>> And like my Sunday night last night was like thinking about our show and the conversation. What a great Sunday night. What a great Monday night.

>> I was I was I was not sure where you were going to go with that.

>> No, no. Like what it was How lucky am I? Seriously. It's amazing. Um final one.

What are you most excited for when you look forward next 5 to 10 years? What are you most excited for? Is it becoming a grandparent? Maybe. I might have just seen my mother become a grandmother.

It's amazing. She's amazing. Is it the health benefits? You know, I'm so excited that AI might be able to solve multiple sclerosis, which my mother has.

That'd be incredible.

>> You never know what tomorrow's going to bring you.

So, it's [snorts] the only way I've been able to do everything I do is not to get too hung up on what's going to happen to me a year from now or 5 years from now, because that's too far.

I think you wake up in the morning given an amazing day and you know, your everything's working around you, your kids are happy, your family's happy, you enjoy what you do, you have good friends. And I was at a different space place the other day earlier this week and they asked me like, you're not a 996 CEO, what do you do? And I said, look, I try to make sure that I can find something to enjoy every day.

Because

I have enough things to worry about. I could really get myself in the wrong headspace by worrying about a lot of things. I run cybersecurity for crying out loud, you know? I live off the fact that somebody's going to hack somebody at some point in time. My phone rings saying, can you help us? Why didn't you spend the money before? But can I help you? Yes, I can help you. So, I think it's a it's a state of mind thing. Can you get your state of mind to be optimistic, positive, and one of gratitude and happiness every day? If you can, it's going to be great. Well, guess what? If the health benefits come, you know, I can start being Benjamin Button, be amazing. If my kids are continuing to be happy and successful, be amazing. My mother, you know, lives for 150 years and she's happy, it'd be amazing. There are so many amazing

things that can happen and perhaps may not happen. So,
let's just focus on tomorrow.

>> And Akash, I so appreciate you being willing to come back for [laughter] a second. I mean, after the first one I suggested it, I was like, there's no chance he's doing this.

>> [laughter]

>> Good thing I have to go back and listen to the first one now.

>> But thank you so much. You've been incredible.

>> Thanks for having me.