

Palo Alto Networks CEO: "AI Found 5 Years of Bugs in 6 Weeks"

31 MIN · YOUTUBE · [HTTPS://WWW.YOUTUBE.COM/WATCH?V=H0BRMV6QCI0](https://www.youtube.com/watch?v=h0bRMv6qCi0)

<https://www.youtube.com/watch?v=h0bRMv6qCi0>

SUMMARY

Palo Alto Networks CEO Nimesh Arora discusses the transformative impact of AI on cybersecurity and business operations. He highlights the significant advancements in AI's ability to identify vulnerabilities in code, the changing landscape of SaaS, and the future of enterprise software as it adapts to new technologies.

- Palo Alto Networks has grown from a \$17 billion to a \$238 billion market cap under Arora's leadership.*
- AI is democratizing intelligence, allowing for more consistent outputs across large teams and improving business operations.*
- The AI tool Mythos has drastically reduced the time needed to identify vulnerabilities in code, finding issues in weeks that would typically take years.*
- The cybersecurity landscape is in a race between defenders and attackers, with increasing risks as AI capabilities grow.*
- Traditional SaaS models may become obsolete as businesses can leverage AI to analyze data without needing intermediary software.*
- The future of enterprise software will likely see a shift towards agent-driven systems that eliminate the need for user interfaces.*
- There is a growing need for infrastructure software to manage and store vast amounts of data as AI becomes integral to business operations.*
- Arora emphasizes the importance of reducing false positive rates in AI models to ensure effectiveness in cybersecurity applications.*

Speaker 1: One of the biggest winners right now,

Speaker 2: the big daddy of the cybersecurity space,

Speaker 1: Palo Alto Networks is an outperformer in the space.

Speaker 3: CEO Nimesh Arora. This might come as news to you, but humans have been writing bad code for a very long time. I spent 10 years at Google and Google Search was democratizing information. If you take that analogy and think about what AI is doing, AI is democratizing intelligence. Money is a way to keep track. Yeah. It's not the goal.

Speaker 2: You've been the CEO of Palo Alto Networks for eight years.

Speaker 3: Coming up in eight years this week.

Speaker 2: Eight years. And I think when you started it was 17 billion dollar market cap, if I remember correctly. And this morning I checked, it's 238 billion, which if you listen to what we said yesterday, now that you passed 100, you're more likely to actually 10x. So the first 10x was actually much, much harder. So you're on your way to a

Speaker 3: trillion dollars from your mouth to God's ears.

Speaker 2: I think you are. Okay, so let's just double click into what you see because you are sort of in a really interesting position to see all of it. You see the birth of AI, maybe you've seen the rise and fall of sas. All the models talk to you.

Speaker 3: You were one of the rise again.

Speaker 2: Right, the rise again. You were one of the first and the few that got access to Mythos. So just let me just push the button. Go Nikesh. Start.

Speaker 3: Well, first of all, thank you for having me here. I think AI is exciting. I think it's exciting to see all the stuff that's gone down the last possibly 24 months. I think Sarah just said it. They were right in anticipating the huge amount of compute that was going to be needed. So all that stuff's going on. But you can see that there's this notion which we talked about briefly last time, that AI is really democratizing intelligence.

Speaker 3: What that means is I have 250 people in marketing. They produce varied forms of output. Now I can get 90% of the output to be consistent across those 250 people. I have 5,000 people who talk to customers. My failure Mode is when 5000 people do different things where people say, I want to talk to Joe because he knows how to solve the problem and Jim doesn't. So now we can get 5,000 people to act almost consistently in their interactions with people on the other side.

Speaker 3: So I think it's going to have a phenomenal impact to how we run businesses, how we operate. It's going to change the entire landscape. Now, in that context, you touched upon mitos, and I know David has been very involved with this. Mitos has shown us that all the bad code that humans have written over the last 50 years can be assessed by AI and shown the vulnerabilities can be shown. We tested for six weeks, and in six weeks, we found what have taken us five to seven years.

Speaker 2: Wow.

Speaker 1: Say that one more time.

Speaker 3: In six weeks, we found vulnerabilities which would have normally taken us five to seven years to

find.

Speaker 2: So, Methos, these are vulnerabilities where?

Speaker 3: Sorry?

Speaker 2: These are vulnerabilities in your own code base or in your customer or in your own code base. Oh, wow.

Speaker 1: So Mythos was not oversold. It was legit.

Speaker 3: The capabilities of AI in being able to assess vulnerabilities in code are real. Not just that. If you put it on ultra mode, which is persistent thinking, so it keeps trying until it gets an answer, you can actually daisy chain vulnerabilities. That is finding a new attack path into your. Into your vulnerabilities. Now, we pride ourselves as a top percentile of companies that test our code because we're in cybersecurity business. If you take that and compound that across all the companies that exist in the world that write their own code, or the 10 million developers write code, this thing is going to find stuff which would have taken us 10 years to find.

Speaker 1: How much did it cost? Like, did you track the token cost? Was it \$100 million? \$10 million?

Speaker 3: No, it was in the low millions. But again, the cost, as Sarah said, the cost curve is going to come down already. OpenAI has got a model which is cheaper and more consistent. Anthropics come out with another model, you buy the hype. It's not hype. It's true. That's capabilities.

Speaker 4: The capabilities are.

Speaker 3: You know that the capabilities are true.

Speaker 4: Yes.

Speaker 3: I mean, you saw IBM announced a project for \$5 billion to fix open source. That's the biggest problem.

Speaker 1: What would have happened if Claude didn't have the restraint and they put it out in the public? Do you think it would have been like a real attack vector and caused chaos in corporations?

Speaker 3: I think we're three months away, if not already, from this being available in the wild.

Speaker 1: Okay, open source.

Speaker 3: Yeah.

Speaker 4: Just three months?

Speaker 3: Yeah.

Speaker 4: Because, I mean, we've been saying that it's roughly six months away before Mythos level capabilities are available in Chinese models, open models, whatever. But you're saying it could be three months?

Speaker 3: Well, look, there's. What is 4.8 is already out, 5.5 is already out. They have similar capabilities. And look, you don't need to crack the hardest code to crack. You just need to find a few vulnerabilities in code that are out there. Just take an old industrial system which is running old code on the edge, you can find that reasonably easily.

Speaker 4: So we're in a race right now between the cyber defenders finding these vulnerabilities and patching them before the cyber attackers do the same thing.

Speaker 3: Yes.

Speaker 4: And how do you feel like we're doing in that race?

Speaker 3: So not as well as we should be doing, which is great for our business, but that's a different story. So look, every company has to go look at their code base and figure out where the vulnerabilities are and fix them. So if you talk to CIOs today, their biggest problem is all the vendors are showing up saying, please patch my piece of boxes, the hardware that you have, please, please patch my code that you have. Because I found vulnerabilities, fix it.

Speaker 3: While the CIS are busy finding their own vulnerabilities to fix their own vulnerabilities, and then this huge thing called open source, which nobody knows quite how to solve.

Speaker 2: So is it fair to say that as model capabilities go up, systemic business risk of large enterprises also goes up

Speaker 3: on the cyber side? Yes. There are antidotes being built by people like us and others where we're going to provide some capability where you don't have to patch everything. But look, Cyrus did something very interesting around harnesses memory and context. The part we don't talk about here is organizations don't have memory and context of everything they do every day. That's why you need to store a lot more data enterprise wide to learn what good looks like and what bad looks like.

Speaker 3: The same problem is in cybersecurity. We need to collect 10 times the data in the enterprise on a cyber perspective to be able to understand how to defend ourselves against the AI attackers.

Speaker 2: Do you think that the traditional companies like the SaaS businesses that have existed in this world, what is their place? As all this knowledge becomes more persistent and stored, what happens to SaaS?

Speaker 3: Well, you see, SaaS is. Bill said SaaS is different pieces. Right. If you're an analytical SaaS company, it's over.

Speaker 2: It's over. What is an analytical SaaS company?

Speaker 3: Somebody that says, I'm going to collect a lot of data for you and analyze it for you. I don't need you to analyze it for me. I can run models against data and analyze them myself. So if you think about, There's a lot of. Every SaaS company has a marketplace you can buy Salesforce Marketplace. What are they saying? You have Salesforce data. I'm a marketplace app. Take me and I'll help you analyze the data. I don't need you.

Speaker 2: You don't need that.

Speaker 3: I can just go run an LLM against the data. So the entire incrementality that has been sold as incremental software modules to all of us doesn't need to be sold to us. Because I'd much rather have LLMs run against the Internet.

Speaker 1: Interesting. You bring this up. We had an instance with a SaaS product with 20 seats. Nobody was logging in and using it. But the data was there.

Speaker 3: Yes.

Speaker 1: So we created like three accounts, got rid of 17, connected it to Slack, connected it to Claude, and now everybody can interface it through a natural language and we've reduced our bill by 90%.

Speaker 3: Well, not just that. What are you going to do next, Jason, is that you're going to take data from different products, put them in one place, the analytics against that, I want my data for my sales reps, my productivity data, my, you know, inventory data from SAP. I want it all one place so I can run analytics against it and say, who's selling a lot? Where do I have less inventory? Let's build inventory in a region where my salespeople are extremely productive.

Speaker 3: To run that query, you'd have to have talk to three different SaaS products. Tomorrow you can put all the data in one place. So that's sort of category one analytical.

Speaker 2: SaaS is dead. Category one analytics dead?

Speaker 3: Yes. In the medium term, you get all these bounces today and tomorrow. These are marginally irrelevant. Infrastructure software undervalued.

Speaker 2: What is infrastructure software?

Speaker 3: Stuff that gives you databases, you collect data into it. Stuff that allows infrastructure to work, whether it's a, you know, database software.

Speaker 2: Databricks, Snowflake.

Speaker 3: Like Databricks, Snowflake, MongoDB, Oracle, all these things.

Speaker 2: You need core storage infrastructure, core data.

Speaker 3: We are going to need 10 times the data stored in enterprise than we have today. Three years. Ten times. Okay, so anything that helps you collect infrastructure data, manage it, you need. I think the category in the middle is called, let's call it system of work or system a record, people call them. Those are deeply embedded in the way businesses work. I have 6,000 salespeople. They know how this works. What's gonna happen is step one, we will take away UI and let agents do the work.

Speaker 3: UI, enterprise software and consumer software. UI is the worst thing we did as technologists.

Speaker 2: You had a couple of examples of this. You told me this story, I don't know if you wanna repeat it, of this one company, they tried to hold you hostage on a license.

Speaker 3: Yes, that was analytical SaaS. So that's okay.

Speaker 2: And you just pointed AI at it and you just.

Speaker 3: Yes, we just got rid of them. That's a different issue. But I mean, think about it. Today we spend our lives having product managers design UI so all humans can interact with data behind the ui.

Speaker 2: Yeah.

Speaker 3: If all, like if you believe agents are going to work and I say I just tell an agent, look, figure out from my sales call, figure out the key points and go post it into, you know, whatever sales tracking system I have with this Oracle or Salesforce. Right. An agent conceptually should be able to do it. Should. We're spending trillion

dollars building these agentic backends. We need these agents to be able to do it. If that happens, UI goes away.

Speaker 3: If UI goes away, I can rewire my system of work. Right. I have sales guy shows up and says I had the sales call do all the paperwork and all the shit that needs to happen in the back end of the company and just I'm done. Right. If I can change the way work happens, which is where you will get true efficiency, where five people become one in a company. All these SaaS software that does system work needs to be re engineered for the next five years.

Speaker 1: And it's also happening passively, which is really interesting. It's looking at email, it's automatically taking the zoom transcript and summary. So the sales system of record is now like you don't even need to input it. It's like I already have the zoom call notes. I have the deck. The deck was made, the sales deck was made by AI. It's just we're all going to be looking at a chat window and just saying, here's what I want.

Speaker 3: Your audit trail is a lot better because humans are not touching your data. It's always been managed by agents. So I think the whole system of work, system of record gets reinvented in the next five years.

Speaker 1: Yeah, there's no data entry. That's an interesting point.

Speaker 3: Yeah.

Speaker 2: Let's talk about national security for a second. I just want to maybe zoom out. So the one side of Mythos as you said, is like the value that it has to you and to your and to enterprises. The Red team version of Mythos is where foreign state actors can essentially create economic havoc inside of a country. As these Models escalate in their capability. What do you think should happen when these models are ready?

Speaker 3: The sad truth is in a year there's a few thousand breaches or attacks that happen. They happen for pretty rudimentary reasons. It's not because somebody cracked a hard to crack thing. It happens because 89% of the attacks happen because credentials get stolen or use

Speaker 2: your name and password.

Speaker 3: That's it.

Speaker 2: My password is password.

Speaker 3: Yeah, I'm sure it is. Do you have dollar sign?

Speaker 2: Dollar sign.

Speaker 3: Fantastic. Well done. See, you're already ahead of Everybody else. So 89% of breaches happen because of simple things. I don't think we need more models to go crack the stuff now. We will need these models can attack critical infrastructure and things we try and protect. From a national security perspective. Yes, we need defenses there. I'm not worried about the national security part being protected because they're very on it. They are the right people. They spend 10% of their budgets on it, on security.

Speaker 3: I'm worried about the small offices across the country where they're using some piece of packet software and you're running a dentist office or a doctor's office. Remember when change healthcare got breached? Every physician's office shut down.

Speaker 2: Shut down in its ransomware because of

Speaker 3: ransomware to change healthcare. That's what the clearing system. That's when United Health act actually have give billions of dollars of credits to the physicians to be able to run their businesses. At that point in time. That's what one should worry about. It's less about the big nuts will get cracked.

Speaker 2: It's less about cracking some PG&E power generation facility. It's more economic chaos.

Speaker 3: Yes.

Speaker 2: And so what, what, what do we do?

Speaker 3: I don't think there's a sort of a silver bullet. I think this will take time. I think this will basically take a while until every system gets upgraded, renewed, fixed over time. I just think it increases the terminal value of the industry. Right, right.

Speaker 2: Do you think that there's a world in which these models become so good that you could see yourself advocating for more nationalism around how they're controlled and how they're managed and where we point them? Or do you think there should be maybe a set of these models that never see the light of day, that only the NSA and other folks could have access to or guys like you?

Speaker 3: I have a slightly differentiated view about models and how they will evolve versus what we heard earlier. From an OpenAI perspective. I think I still believe models are going to become a utility layer. You'll be able to buy intelligence on the fly. Or you can say, I don't need 180 IQ person to go do this task. Give me 120 IQ and I need a 250 IQ. Do this task. I'll pay \$10 for this. For this I'll pay \$0.01.

Speaker 3: So I don't know there's a one size fits all. Give you the most up to date model to answer my customer call saying, sorry sir, I have no idea how to solve your problem. So I think models will get differentiated from utilitarian perspective. So if you look at already what's happening in the market, right, the profit pools are in applications, not in models. Sarah talked about Codex running away. She didn't say OpenAI is running away.

She says Codex is running away.

Speaker 3: Just say just the way I'm sure Dario says Claude Code is running away. So you're seeing that they're attacking profit pools. They're attacking profit pools because that's where the money is going to come from. The profit pools are an application that companies can use. The profit pools are not in model usage by companies because most companies have no idea how.

Speaker 1: How do you look at these companies in a way, OpenAI Anthropic as the new Microsoft Office coming in and doing all applications, all productivity software for organizations.

Speaker 3: No, I see there's going to be application companies which are going to arbitrage between models and solve your business problem.

Speaker 1: So you still think they won't go to the application? Larry, because this is a big debate. Should you engage with open AI and train their systems to then take your business from you? And Anthropic keeps releasing their legal model, their accounting model. And it does feel like in order for them to hit their revenue numbers, they might need to do what Microsoft did, which is release the Office product on top of the operating system.

Speaker 3: If I'm a company, I don't want to write every piece of software myself. I want my HR system software, which is Agentic enabled, enabled AI enabled to be delivered by some application company. Could be a new AI application company. I want my sales management system built by the new agentic AI salesforce of the world, whether it's Salesforce or somebody else. So I want applications. Now. What Sarah said is the profit pools are in the application layer.

Speaker 3: That's why they want to be the application layer. So I think we're still waiting for that layer of companies to be invented or created where applications will sit. Because 50,000 companies need the same application. Why would I build it myself? It's Highly inefficient. It's silly for me to use OpenAI directly and rewrite my entire sales system because I'm smart. I'm not. I want somebody to do it for me. So I think that layer of companies is still not fully formed.

Speaker 3: We're still going to be waiting for.

Speaker 2: You want a control plane, a harness

Speaker 3: and then that's right. They will build the harnesses and the memory into those application layers. Now the question is, how big is the application layer? Is it one application, is one, you know, enterprise application that does everything or is it specialized application?

Speaker 1: You did it and you kicked out this software vendor. You did it because they were being abusive in

pricing.

Speaker 3: So we still use this different vendor.

Speaker 1: What's that?

Speaker 3: We swapped out for different meta. We just took more control.

Speaker 1: Love it. So it really is a pricing issue and that's why the SaaS apocalypse in some ways makes sense. They're not having pricing power because you could say I'll just put 10 developers on this and I'll save \$10 million.

Speaker 3: Yes. I think the part back to what Chamath said about the regulation or whether you want to regulate these higher powered models, the question is at some point in time when these newer models, which are even more powerful, get built, they will come at a different price point and they might have to go to a certain vetting process to understand what their capabilities are. But I think we're in a global race. I don't think holding back our models for three to six months is going to help us any somebody else is going to put them out in open source.

Speaker 3: I was shocked to hear when I was talking to the CEO of one of these model companies. He says the entire weights of their most recent model can fit on USB stick.

Speaker 2: Say that again. The entire weights.

Speaker 3: Entire model weights of their newest model fits on a USB stick. That's the ip. That's incredible because all the data can be distilled in under 24 to 48 hours and model comes out.

Speaker 1: I'm curious.

Speaker 3: So that's the ip. So are you telling me that you know we're going to hold on to that for six months?

Speaker 2: Right.

Speaker 1: We have a debate about how difficult it is to make a frontier model. Some companies are starting to think about making frontier models using their data advantage to build their own.

Speaker 3: Have you thought about that?

Speaker 1: Palo Alto? Because it does seem like you have proprietary knowledge on how security works. Could

you build your own large language model or a vsml? A small language model that would give you some advantage.

Speaker 3: Here's the Part nobody talks about is the false positive rates on the models. What is the false positive rate on 4.8 and 5.5?

Speaker 1: No idea.

Speaker 3: You guys don't talk about. You should. The false positive rate on mythos was 30%.

Speaker 2: Oh, wow.

Speaker 3: Right.

Speaker 1: So it thought it found something, but it hadn't.

Speaker 3: Yes. So the problem is it's great for attack, it's horrible for defense. It finds. 30% of the time it finds something. I found a problem. And you say, let's plug the hole. Wait, there wasn't a hole there in the first place.

Speaker 1: No. Missile inbound, right?

Speaker 2: Yeah.

Speaker 3: So now the same problem applies in enterprise. If you use a model without the right harnesses, the right training, you could be running into 10, 20% false positive rates. Let's use the model to pay, I don't know, insurance claims. Yeah. Oh, great. 10%, 20% false positive. I just lost money.

Speaker 1: The sycophantic nature of these is ridiculous.

Speaker 3: So the problem is not who wants the newest model. The problem is how do you take that model with 20% or 10% false positive and make it 0.01% false positive? In my business, I want 0% without

Speaker 5: losing the false negative. Sorry. Without losing the negative. The false negative.

Speaker 3: But it's like saying, hey, let's take the new self driving car. Mercedes is going to use Opus 4.8 and you can just sit in the car, it's going to drive you. I'm not putting my kids in that car with a 10% false positive rate. Are you?

Speaker 5: So there's a lot of work that

Speaker 3: happens post the model, which needs to happen to make this thing useful and effective in the business context.

Speaker 2: Let me slightly pivot for a second. You were for a very long time the chief business officer at Google. You were the president of SoftBank. Now you're the CEO of Palo Alto Networks. So let's play armchair CEO. Armchair CEO.

Speaker 3: I'm still bristling from David Friedberg trying to create a distinction between founders CEOs and non founder CEOs. Just saying. Just saying, David.

Speaker 5: By the way, false positives, false negatives too.

Speaker 2: Give us what you would keep, what you would change, and what you like about the following companies.

Speaker 3: This is going to get recorded and put out there. I don't like that.

Speaker 2: We're asking you a question.

Speaker 5: Don't.

Speaker 3: Hazing people and grossing people.

Speaker 2: Are you ready?

Speaker 3: Yeah, sure.

Speaker 2: Okay. Why you keep what you change, what you like, what you don't like.

Speaker 3: Uber in a world of it. Dude, I can't talk about my. I'm the board of Uber. I'm not going to talk about it.

Speaker 2: I didn't know that. Sorry.

Speaker 3: He's the CEO. He's a great guy. Okay.

Speaker 2: Waymo.

Speaker 3: Trying to get me fired.

Speaker 2: Waymo.

Speaker 3: What do I like about the cars work? It's amazing. They should have more in many more cities around the world faster. I would say that to take it.

Speaker 2: I think she knows Google writ large.

Speaker 3: I think Google's underrated. I think it's going to be the first \$10 trillion company in our lifetime. I think they have all the assets that are needed to make this successful. I think people underestimate. You can be a model company. You still need to have a salesforce that convinces customers to go out there and embrace these models and buy them. If you think about it, three hyperscalers have the biggest number of salespeople out there. So they should.

Speaker 2: Part of why they're a little bit undervalued is just the conglomerate nature is hard to understand.

Speaker 3: I don't know. You guys are smart of that stuff. I'm just a hired hand CEO.

Speaker 2: I didn't say that. Rebrick said that.

Speaker 3: Let's just be clear. I know, I know.

Speaker 5: I was providing a thesis on recovery out of the SaaS apocalypse.

Speaker 3: Okay.

Speaker 5: Just to be clear, there's a way to segment that basketball and you're not in that basket either.

Speaker 3: You're making a distinction about how people who are founders, CEOs have. Have the right to take more risk and are allowed to take more risk.

Speaker 5: I was saying that.

Speaker 3: And I think.

Speaker 5: And I. And I think you. You provide a unique counterpoint to that. And there's not a lot of the

Speaker 3: false positive in your mind.

Speaker 5: I think the same was true of Jeff Weiner. And I think that there's a few other really great CEOs, but they are like Neo in the Matrix type anomalies. And I think you're one of those people. And there's a very rare kind of personality profile of someone that's willing to take risk and take ownership of something that wasn't theirs in the first place. And they make it theirs. And it's an extraordinarily unique trait. Far more unique actually than being a scalable founder.

Speaker 1: It's forgiven.

Speaker 3: You're forgiven.

Speaker 2: Yeah. Good save.

Speaker 1: Credible save.

Speaker 3: Let's go back to ARM three.

Speaker 1: Wow, that was incredible.

Speaker 3: Open the eyes.

Speaker 2: He did really well.

Speaker 1: Sycophantic than chatgpt.

Speaker 5: He's like, actually, I'm actually the best.

Speaker 2: Let's go back to armchair.

Speaker 3: I'm liking this. You guys open AI more often. Yes. They need to sell faster.

Speaker 2: OpenAI.

Speaker 3: They should sell faster, right?

Speaker 2: They should sell faster.

Speaker 3: I mean, I mean you said it, didn't you just say it when Sarah was here? That Anthropic seems to have improved their ARR much faster than OpenAI.

Speaker 1: I mean that's just statistics.

Speaker 4: They kind of went all in on enterprise and coding specifically.

Speaker 3: I think that's like the, the conversation right now is it's a race to take over the profit pools. If you are going to need tens and tens of billions of dollars every year to get what is that? One gigawatt is 10 billion revenue. What does it cost to build you? So what are the most exciting cost? 50. So this is a great deal.

Speaker 4: So what are the most exciting profit pools then? So we got coding. That's been the breakout application over the past year.

Speaker 3: It's massive.

Speaker 4: You've got infrastructure like you said, the new databases, I think cybersecurity, it's clearly one of them because of the threats and patching cycles. So much more dynamic.

Speaker 3: There's a slight difference. Yes. As you can see, these models are trying to be the enablers of better cybersecurity, which is good because all of us need to use them to test. And you're probably going to see if you saw Anthropic has already made their cyber capable model available generally so that everyone can use it. And OpenAI has got one. I'm sure Google has one too. Cause they understand this is a place where CISOs or chief security officers want to use it to test the code.

Speaker 3: So this is another profit pool. I think we haven't seen the onslaught against the application software companies yet. I mean there's tens and tens of billions of dollars in application software which is waiting to get reinvented. As we talked about. I think eventually you'll see these people saying what if I took this 40, 50, \$100 billion TAM down, I can build a whole brand new backbone with a gen tech AI and it'd be so differentiated that it'll cause customers to move.

Speaker 1: We are seeing it as a playbook in the accelerators now in the year zero and year one companies, people are coming to us with the pitch. This is \$1,000 a seat per year, \$500 a month seat, SAS software, we can do it for less. We're going to charge them based on consumption. We're going to take 80, 90% of the cost out as doing with 89

Speaker 3: fastest places to make revenue in enterprise are replacement TAMS. If you replace something I Already have a budget. It's easy. I take something bad or replace with something better. I get money. So replacement tams are beautiful. If you can replace an industry, replace the profit pool is great. The second place is consumer revenue. It's a lot easier to get five bucks per user on a consumer side, Netflix. So that's where I

mean look at it. I think we collectively probably pay more on subscriptions per month than we ever did historically.

Speaker 3: And you thought your cable bill was high.

Speaker 1: Yeah.

Speaker 2: Do you think that you're going to end up building more or less hardware in the future?

Speaker 3: If you have to guess, Hardware, even today is the cheapest way to manage low latency, high throughput bits. You still need a data center. What's the data center? It's just managing high throughput, low latency bits. That's why if you look, financial services is the most reluctant industry to go to the cloud because you increase latency. If you increase latency, you reduce profit. So if you look at every of the largest financial services companies, that is Goldman or JP Morgan, Morgan Stanley, street, these guys, they're doing hardware.

Speaker 2: Yeah.

Speaker 3: Try to get them to run their business on the cloud. They can't because they will have higher latency, they will lose money.

Speaker 2: Right.

Speaker 3: So hardware still be made. I mean I remember when I used to buy Silver Lake and I had heard Dell was done, nobody wanted hardware. I think Dell might be back to like a 3,400 billion dollars market cap. So hardware is still going to be around, we're going to need it. It's the fastest way to move.

Speaker 5: Are hardware development cycles changing because of AI? Are you seeing a lot of generative design stuff moving in silica that historically was manual and long cycle?

Speaker 3: Yeah, but the long pole in the tent is never designed. Right. Long pole in the tent is production today. Production. You can't get a box produced because every, every piece of hardware componentry is backordered. Everything's expensive and every factory in the world is backordered because we're trying to build all these GPUs based, you know, chip cards for every data center of the world.

Speaker 5: So you think the US is equipped to fill that supply chain need, can we do that here or do you think we've done 10 years with a commit, with a, with a firm top down commitment?

Speaker 3: Well, I mean the good news is that I think the hardware industries is seeing a bonanza of a lifetime. And generally when you see a bonanza lifetime, you can go commit 10, 20, 50, \$100 billion. I mean I've seen a

CEO on television committing a hundred billion dollar plan to go build more memory. So that's good. That means they have the money to go put the money in the ground literally to go build these things for the future.

Speaker 3: So I think that gets us more

Speaker 5: certain that the, I think the tax incentive has a big, has a lot to do with that. The accelerated depreciation on the, the capex, you get 100% write off in the first year right under the, under the,

Speaker 2: Just a, just a final question as we wrap up you. Over the last eight years you've grown organically, very aggressively, but you've also been pretty acquisitive. You'll, you know, you'll take shots and they've generally worked. So you have a ton of permission in the market. When you hear what Bill Hackman said about how there's this kind of overbeaten companies, there's a few that get celebrated, that's a ripe pool for you to pick from. But some of that would require you to go maybe a little horizontally, far afield.

Speaker 2: Some would say how do you maintain the discipline? Or do you see yourself at some point considering things that are not nearly so much right down the middle of cyber?

Speaker 3: So I tell you what, until about a year and a half ago, we used to buy product companies and throw them into our go to market engine. We could rewire their backend so they can work better with our go to market engine. So for me, if I'm selling \$10 million to a customer next time I go two years later, if I can sell them 20, it's the most efficient way for me to amortize my go to market spend.

Speaker 3: So that was the model we played that we ran that playbook to north of 150 billion. Then we got to a point where it says, oh, we see an inflection arriving in identity. It's going to be important from an agentic perspective, security perspective. So you bought a \$25 billion company which we closed three months ago. Now it's actually a very different opportunity has presented itself and the different opportunity sort of goes like this. If you can be the best at leveraging AI to run the most efficient enterprise business in the world, your operating margin can be far in excess of the industry.

Speaker 3: And if you can, if you can

Speaker 2: crack that code gross and net, you're saying GROSS in the 90s, net in the 40s.

Speaker 3: Yeah. If you can crack that code, then it doesn't matter what you buy. I think the problem right now is the execution problem. Most subscale companies cannot afford to go optimize their company and run it better. So if we can run our company much better than everybody else and have a higher operating margin, then the street will say, fine. If you take something at a 20% margin, make it A.

Speaker 2: Your first M and A was really tough. No, like they were pretty skeptical and then you kind of shoved it in their face.

Speaker 3: They're pretty skeptical when they found a guy who didn't know cybersecurity, didn't know enterprise show up, who worked at Google. And the track record of people leaving Google and being successful out of Google is still varied.

Speaker 2: So basically you're saying the menu is open.

Speaker 3: I think we need the next six to 12 months to figure out how this AI settles down and how can we use that effectively in enterprises. I think if you think about it, the people keep hoping that less people we need to run companies. I actually have a counter view. I think we have more people at Palo Alto on the technology side than we've ever had before. Because I think AI is causing everything to ask for a transformation.

Speaker 3: So I have more technical people today than I would have had if AI didn't exist.

Speaker 2: Ladies and gentlemen, CEO of Palo Alto Networks, Nikesh Arora.

Speaker 3: Thank you, guys.

Speaker 1: Thank you, sir.