Corrected Transcript

18-Nov-2026

Palo Alto Networks, Inc. (PANW)

Q1 2027 Earnings Call

Copyright © 2001-2026 FactSet CallStreet, LLC

Total Pages: 22

CORPORATE PARTICIPANTS

Hamza Fodderwala
Senior Vice President-Investor Relations, Corporate Development and Strategic Finance, Palo Alto Networks, Inc.

Dipak Golechha
Chief Financial Officer, Palo Alto Networks, Inc.

Lee Klarich
Chief Product and Technology Officer & Director, Palo Alto Networks, Inc.

Nikesh Arora
Chairman & Chief Executive Officer, Palo Alto Networks, Inc.

OTHER PARTICIPANTS

Saket Kalia
Analyst, Barclays Capital, Inc.

Brian Essex
Analyst, JPMorgan Securities LLC

Meta A. Marshall
Analyst, Morgan Stanley & Co. LLC

Patrick Colville
Analyst, Scotiabank

Shaul Eyal
Analyst, TD Cowen

Fatima Boolani
Analyst, Citigroup Global Markets, Inc.

Jonathan Ho
Analyst, William Blair & Co. LLC

Michael Turrin
Analyst, Wells Fargo Securities LLC

Joshua Tilton
Analyst, Wolfe Research LLC

Gabriela Borges
Analyst, Goldman Sachs & Co. LLC

John DiFucci
Analyst, Guggenheim Securities LLC

MANAGEMENT DISCUSSION SECTION

Copyright © 2001-2026 FactSet CallStreet, LLC

Hamza Fodderwala

Senior Vice President-Investor Relations, Corporate Development and Strategic Finance, Palo Alto Networks, Inc.

Good day, everyone, and welcome to Palo Alto Networks Fiscal First Quarter 2027 Earnings Conference Call. I am Hamza Fodderwala, Senior Vice President and head of Investor Relations. Please note that this call is being recorded today, Wednesday, November 18, 2026 at 1:30 PM Pacific Time.

With me on today's call to discuss our fiscal first quarter results are Nikesh Arora, our Chairman and Chief Executive Officer; and Dipak Golechha, our Chief Financial Officer. Following our prepared remarks, Lee Klarich, our Chief Product and Technology Officer, will join us for the question-and-answer portion.

You can find the press release and other information to supplement today's discussion on our website at investors.paloaltonetworks.com. While there, please click on the link for Quarterly Results to find the Q1 2027 supplemental financial information and Q1 2027 earnings presentation.

During the course of today's call, we'll be making forward-looking statements and projections regarding the company's business operations and financial performance as well as the company's recent acquisitions. These statements made today are subject to a number of risks and uncertainties that could cause our actual results to differ from these forward-looking statements. Please review our press release and recent SEC filings for a description of these risks and uncertainties. We assume no obligation to update any forward-looking statements made in today's presentation.

This presentation also contains non-GAAP financial measures and key metrics relating to the company's past and expected future performance. Non-GAAP financial measures should not be considered a substitute for financial measures prepared in accordance with GAAP. The most directly comparable GAAP financial metrics and reconciliations are in the press release and the appendix of the investor presentation. Unless specifically noted otherwise, all results and comparisons are on a fiscal year-over-year basis.

With that, I will now turn the call over to Nikesh.

Nikesh Arora

Chairman & Chief Executive Officer, Palo Alto Networks, Inc.

Thank you, Hamza. Good afternoon, everyone. As you can all see, we had a strong start to our first quarter of the fiscal year. We beat guidance across every metric. NGS ARR was up 66%, and we delivered one of the largest NGS ARR beats in our history. Total revenue was up 39% and RPO grew 35%, both of which accelerated versus the prior quarter.

Delivering that calendar result in a seasonally slower Q1 speaks to the strength of the demand function. As AI accelerate cyber attacks, enterprises are moving with urgency to replace disjointed tools with unified platforms. That momentum gives us confidence to materially raise our guidance for the year.

Top line strength also translated directly to the bottom line. In Q1, we once again exceeded 30% operating margin, setting a new record for operating profitability while integrating large acquisitions. Most encouraging is how demand was broad-based across our platforms, all of which delivered accelerating revenue growth.

Copyright © 2001-2026 FactSet CallStreet, LLC

Starting with our largest business, Network & AI Security, which grew 20% with strength across SASE, software, and hardware firewalls. Customers are moving beyond point products and recognizing that running fragmented perimeter defense across multiple vendors leads critical gaps in visibility and enforcement.

We're also seeing a fundamental shift in what traverses the network. Agentic traffic on SASE is up approximately nine times year-over-year as autonomous agents continuously connect to applications, APIs, and enterprise data, all of which require real-time inspection. That surge in traffic reflects enterprises moving AI into production and securing this new surface area has propelled Prisma AIRS to nearly $170 million in ARR in five quarters since it was generally available, making it the fastest growing product in the history of Palo Alto Networks. Cortex grew over 50% in Q1, led by XSIAM and observability.

As AI-driven attacks move at machine speed, BAND-AID reactive solutions across multiple data lakes are not a plan. We need solutions that run real-time analytics and observability to dynamic correlation and are able to deliver real-time correlation and verdicts. XSIAM now in just over 20 petabytes of daily telemetry, leveraging AI to reduce response times from days to under 10 minutes.

Turning to Idira, our identity security platform revenue here grew 23% in Q1 on a pro forma basis. We continue to execute on our integration plan since closing the CyberArk acquisition early this year. We added more than 100 net new logos from existing Palo Alto Networks' installed base in Q1, our best cross-sell quarter-to-date between Idira and Palo Alto Networks.

The rise of AI agents introduces an entirely new class of non-human identities that must be secured in real-time. Maintaining effective oversight requires complete visibility and the capability to instantly revoke access when suspicious activity occurs. To address this exact challenge, in Q1, we launched general availability of Idira's agentic IdP, integrated directly into our Prisma AIRS AI gateway, enabling real-time enforcement against agentic permissions.

Across all of our three platforms, customers are solving the same core problem, compressing the time between spotting exposure and shutting it down. Point products simply cannot keep up when attacks move at machine speed. Enterprises are responding by consolidating their security architecture onto unified platform, an ongoing shift that continues to provide momentum for platformization, yes, that word platformization.

In Q1, net new platformizations hit a record of approximately 170, which was more than the last two Q1s combined. Platformized customers now represent two-thirds of our Next-Generation Security ARR and continue to expand a robust rate with net retention of up 120%.

Our largest deals detail the platformization strength we saw in the quarter. In Q1, we closed an $83 million deal with a global financial institution. This customer platformized our network security, adopting frontier virtual patching and expanded identity security, signing the largest ever deal in CyberArk history. A global consulting firm signed a $46 million transaction, which included our largest ever Prisma AIRS deal as they move AI into production across their organization.

Finally, we signed a $20 million deal with a global IT services firm to help unify their security stack and reduce their point product sprawl and modernize their defenses. This customer purchased Prisma AIRS and SASE to secure AI deployments and access across their global workforce.

Copyright © 2001-2026 FactSet CallStreet, LLC

Customers are making long-term architectural bets with us, and our responsibility grows with every commitment. They expect us to keep sharpening the products they rely on today while staying ahead of what comes next.

Six months ago, when cyber-capable models like Mythos were first released, we predicted that an equally powerful offensive AI capability would rapidly diffuse through open-source within three to six months, prediction that has now materialized.

I want to be clear. There are two distinct parallel effects that are driving this environment. The first effect is a massive surge in vulnerabilities and patches being deployed. We're seeing critical and severe vulnerabilities discovered nearly seven times the monthly average of last year. These discoveries are happening at a speed and intensity which is hard for IT teams to cope with. We believe many of the major disruptions the last few weeks, including outages across major enterprises and governments, were partially due to unintended consequences of patching and configuration changes as IT teams simply cannot patch these vulnerabilities fast enough without compromising their production environments.

The second effect is the rise of AI-assisted cyber attacks. Bad actors are using open-weight models and new agentic frameworks to find and exploit vulnerabilities at unprecedented speed and scale. New agentic frameworks like ARTEX are enabling bad actors to turn simple instructions into autonomous, multi-stage attacks, allowing a handful of operators to compromise dozens of targets simultaneously.

We saw an example of that just days ago when our Unit 42 team published research on a campaign that involved a small team of foreign-linked hackers that were able to infiltrate 19 organizations at the same time. As we look ahead, we anticipate an elevated threat environment, but we believe enterprises have multiple ways to adapt to this new paradigm.

First, the significant volume of patches will continue to be required for at least the next 18 months, carrying ongoing risks of collateral damage. We expect the collateral damage will decrease over time as customers get smarter about understanding exactly what to patch and when to patch.

Second, the need for virtual patching will grow across all industries and software categories, allowing organizations to securely and comfortably shield their infrastructure in line before applying their patch. In Q1, we saw strong demand for our advanced virtual patching service, and we see a large opportunity to deliver this to our nearly 70,000 network security customers.

Third, the reality for every enterprise is that there is no silver bullet for preventing an attack, and organizations can no longer assume it won't happen to them. When an attack hits, because it will, defenders have to be ready to detect and remediate in record time rather than relying on reactive BAND-AIDs. Therefore, we believe the third part of the antidote is most important, which is modernizing the north of $1 trillion of legacy security and IT infrastructure onto real-time platforms.

Despite the heightened threat environment, demand for AI remains insatiable. Every drop in inference cost unlocks thousands of new genetic workflows, putting industry token revenue on a trajectory to surpass $1 trillion by 2029.

While the market is focused on AI trade, we are focused on AI inflection, the shift from building AI infrastructure to real enterprise consumption. Until recently, most AI activity has been concentrated in consumer use cases and horizontal applications like coding where there's virtually no security around these tokens.

Copyright © 2001-2026 FactSet CallStreet, LLC

Going forward, we believe the currently deployed AI infrastructure will also begin to show vulnerabilities, and we've seen recent examples of these exploits among AI infrastructure providers. As coding becomes more agentic, these lack of security controls will bring agentic risks into enterprise environments. It is imperative that we ensure AI adoption is in lockstep with AI security.

We're already helping our customers close the gap with Prisma AIRS, reaching near $170 million in ARR across north of 1,000 customers within the last five quarters. Today, Prisma AIRS delivers comprehensive AI security across five pillars, giving enterprises a single platform to see, assess, and control their entire AI footprint.

It starts with AI discovery and AI risk assessment. First, can you find every AI artifact in the enterprise? Prisma AIRS then discovers models and autonomous agents running anywhere then continuously evaluates the risk across your supply chain, security posture, and automated AI red teaming.

From there, we move to real-time enforcement through AI protection, AI governance. Can you consolidate all that traffic into a single enforcement point? With runtime security, identity security, and our AI gateway working together, customers inspect prompts, block sensitive data leaks, secure agent identities, and enforce policy across every interaction in real-time.

This is the industry's first AI security platform and tying it all together is AI observability, capturing and observing every prompt, response, and log so organizations maintain complete accountability at AI scale. While Prisma AIRS encompasses important pillars of AI security today, we expect that as models continue to get smarter, they will obviate static human controls. Because of this, AI security models will need to be equipped with advanced reasoning capabilities that are on par with the frontier.

Trusting LLMs will require that we can independently verify, explain, and audit AI implementations. We believe both trust and control be essential to driving enterprise adoption of AI.

We're already seeing the desire for sovereignty and control compel enterprises to adopt open-weight and opensource models, and early signs that shift are clear with open-weight models growing from under 5% of token traffic to more than 60% over the past year. When organizations put AI at the center of their operations, they want to retain control of their proprietary data, institutional judgment, and competitive edge inside their own four walls.

We expect a future with a serving opportunity in the hundreds of billions of dollars where enterprise applications are driven by open-weight and open-source adoption. Trust coupled with ability to control AI implementation where it's both trainable and available on-prem or in a sovereign stack will accelerate enterprise adoption of AI. That means that these models must be delivered securely and where required in a sovereign fashion, with controls strictly in the hands of the enterprise.

When we looked at training custom open-weight models ourselves, we looked for partners across the entire AI ecosystem, looking for rare engineering rigor and a proven ability to address enterprise challenges. We landed on Applied Compute because no other team matched their technical pedigree in post-training open-weight models. We're now bringing a world-class team of frontier AI researchers in-house and expect to reach 100 researchers by year-end.

Copyright © 2001-2026 FactSet CallStreet, LLC

Every enterprise sits on decades of proprietary data, internal playbooks, and institutional judgment that cannot be exposed to public models. Until now, they have lacked a trusted partner to post-train, run, and supervise openweight models safely on their own.

We believe Applied Compute is the right partner for us to address this opportunity. Team has already seen great traction with large enterprises and AI natives alike, achieving north of $100 million in annualized revenue with the last 18 months of founding, putting them in rarefied air.

Applied Compute will allow us to execute on two major fronts. First, we will embed AI across all of our security capabilities and are in the process of training north of 35 specialized use cases and making them available to our customers to fight AI with AI.

Secondly, we will help train and serve secure sovereign implementations of open-weight models to the enterprise. Open-weight model implementations will allow us to deliver secure open- weight models to our customers to enable adoption of secure enterprise – secure AI in the enterprise.

To support our efforts, we plan to build out a diversified global inference capability of north of 100-plus megawatts over the next three to five years. This will be supported by a [ph] modest (00:14:00) CapEx commitment of at least – of less than $3 billion through fiscal 2030, representing approximately half a year of our annual free cash flow. By pairing their platform with our global go-to-market reach and the trust we have earned with CIOs worldwide, we expect Applied Compute to continue scaling rapidly and add over north of $1 billion of ARR to our NextGeneration Security ARR by fiscal 2029.

In closing, Q1 was a defining inflection point for cybersecurity. AI-driven attacks have stripped away the luxury of incremental change, proving that fragmented, decade-old tools cannot stop autonomous threats and creating an urgent mandate to modernize $1 trillion of legacy infrastructure.

At the same time, to unleash AI adoption across the enterprise, AI security must evolve from visibility and enforcement to ensuring trust in the model itself. Organizations will be reluctant to put mission-critical workflows on AI until they can rely on how those models behave in production.

Lastly, we see a massive opportunity in enabling secure adoption of open-weight models as sovereignty and control become increasingly important pillars in the enterprise. By bringing that foundation to our customers alongside our core security platforms, we sit at the center of both their security modernization and the AI transformation for years to come.

With that, I'll turn it over to Dipak to cover financial details in more depth and share our updated outlook.

Dipak Golechha

Chief Financial Officer, Palo Alto Networks, Inc.

Thank you, Nikesh, and good afternoon, everyone. We delivered a record Q1 result with accelerating growth on all of our key top line metrics. RPO grew 35% year-over-year to $21.0 billion. Current RPO was up 36% yearover-year to $9.4 billion. Contract durations remained flat year-over-year.

Copyright © 2001-2026 FactSet CallStreet, LLC

NGS ARR reached $9.73 billion, up 66% year-over-year and net new NGS ARR more than doubled year-overyear. We significantly exceeded our guidance to deliver one of the highest quarterly NGS ARR beats in our history.

Total revenue was up 39% year-over-year, reaching $3.44 billion in Q1. As we announced last quarter, we're now reporting revenue across our three segments: Network & AI Security, Cortex, and Idira. Q1 revenue performance was truly broad-based with all three of our segments delivering growth acceleration.

On Network & AI Security, that delivered growth acceleration at scale in Q1, coming off a very strong Q4. Revenue grew 20% year-over-year to $2.3 billion. We saw strong growth across all form factors in network security. Starting with SASE, ARR grew 30% year-over-year, and we added over $100 million in net new NGS ARR. Demand was particularly strong for SD-WAN and secure browser as customers continue to favor our comprehensive SASE platform over point product vendors.

In firewall, we saw strength across hardware and software, with trailing 12-month bookings up 28% year-overyear. We are continuing to gain share as the desire for unified visibility and prevention is driving vendor consolidation within our customer base.

Software firewall ARR accelerated for the third consecutive quarter and grew 34% year-over-year. This was largely driven by Prisma AIRS, which continues to be our fastest growing product, reaching nearly $170 million in ARR.

Turning to Cortex, our security operations and observability platform, Q1 revenue grew 53% year-over-year and 40% year-over-year on a pro forma basis. XSIAM ARR grew over 65% year-over-year to approximately $800 million as customers prioritize their SOC modernization projects to combat the heightened threat environment driven by AI.

Furthermore, our observability platform, recently relaunched as Cortex XCOR, continued its strong growth. ARR was up approximately four times year-over-year, and as we commented last quarter, that includes an approximately nine figures of net new ARR from a large frontier lab customer migrating to Chronosphere from an incumbent provider.

In Idira, our identity security platform, Q1 revenue grew 23% year-over-year on a pro forma basis. We continue to see strong go-to-market collaboration between our core sales team and Idira. In Q1, this resulted in over 100 net new logo wins for Idira, which came from our existing core installed base, representing our largest number of cross-sell wins since the acquisition closed.

We're also seeing synergies – strong synergies on the bottom line. As a reminder, when we announced the acquisition, we targeted Idira operating margin converging with our core business within 18 to 24 months postclose, and are on track to do that in approximately half the time.

From a geographic perspective, we saw broad-based strength. Americas grew 42% year-over-year, EMEA grew 40% year-over-year, and JPAC grew 26% year-over-year. Complementing our record top line performance in Q1 was our continued focus on profitable growth. This discipline yielded continued operating leverage year-over-year, reinforcing our confidence in our margin guidance targets for the full year.

Copyright © 2001-2026 FactSet CallStreet, LLC

Our scale, our diversified business model helps us mitigate gross margin headwinds, and we continue to deliver gross margin in the mid-70% range. In Q1, our total gross margin was 75.2%, down 170 basis points year-overyear. As we noted last quarter, we continue to actively manage multiple factors impacting our gross margin over the near to medium term.

First, we are seeing industry-wide elevated component costs within our hardware business, which we're proactively leveraging our scale and strategic supplier relationships. While hardware represents approximately 10% of our total revenue, we expect the elevated component cost to impact our gross margins throughout the course of the fiscal year.

Second, our gross margins were impacted by a greater mix of higher growth SaaS offerings that are earlier in their maturity curve. This is particularly true for the Cortex segment where we are seeing our XSIAM and observability business grow rapidly as we focus on taking share.

Additionally, our ongoing Prisma to Cortex cloud migration remains a temporary drag on our segment margins. Given the migration and shifting product mix, we expect our gross margin in Cortex to remain below our overall company average over the medium term before returning to 70%-plus gross margins longer term.

From an operating margin standpoint, we continue to offset higher COGS with improving OpEx leverage as we scale efficiency. In Q1, our non-GAAP operating margin was 30.7%, a 50-basis point increase year-over-year despite higher costs related to M&A and gross margin impacts.

We are also driving a number of AI efficiency plays targeted in areas across our business. As an example, this includes our AI quoting efforts to accelerate our time to close in an effort to drive higher sales productivity. In addition, we're driving AI efficiencies in procurement and have already identified significant savings through our vendor rationalization efforts.

Turning to the bottom line, non-GAAP diluted EPS was $1.05, up 13% year-over-year, beating guidance by $0.08 whilst EPS – GAAP EPS was $0.57, up 21% year-over-year. Our improving profitability contributed to strong Q1 cash generation with adjusted free cash flow of $1.9 billion, representing a 56% margin.

Regarding capital allocation, we had no share repurchases in the quarter, and our approach continues to be opportunistic. We continue to have authorization of up to $1 billion in share repurchase until December 2027. Our overall performance is the direct result of our continued focus on innovation. Through organic investments and targeted M&A, we strive to continually stay ahead of technology inflections and anticipate our customers' needs years in advance.

As Nikesh noted earlier, we're excited to announce our intent to acquire Applied Compute to position ourselves as the pioneer in delivering secure AI. The total consideration is $4.5 billion paid entirely in stock, and we expect to close in our third quarter of fiscal 2027. The acquisition includes approximately $500 million in cash and prepaid compute commitments owned by Applied Compute.

The company has already achieved impressive growth, reaching $100 million in annualized revenue within 18 months of their founding. This transaction will be top line growth accretive from day one, and we expect to scale the business to add $1 billion in NGS ARR by fiscal year 2029.

Copyright © 2001-2026 FactSet CallStreet, LLC

Longer term, we believe this is our next multi-billion dollar ARR business and positions us to be at the frontier of AI security. While we expect our compute needs will grow to deliver secure post-training and inference for our customers, we intend to remain capital-light and we'll maintain strong free cash flow margins through and beyond the CapEx cycle.

To meet our compute needs, we will continue to source from our cloud partners and over time add dedicated capacity of our own. That investment will mostly go to GPUs and servers with a smaller share in the facilities that house them. We intend to build our capacity in stages against demand that we have visibility into.

In total, we anticipate our cumulative CapEx needs through fiscal 2030 to be less than $3 billion, which is approximately half of our expected free cash flow this year. Assuming a Q3 2027 close, we expect the transaction to lower our operating margin by up to 100 basis points in fiscal 2027 and up to 200 basis points in fiscal 2028 before returning to our current 29% to 30% operating margin in 2029.

Lastly, from a free cash flow perspective, we expect to maintain 38%-plus margins excluding data center CapEx in fiscal 2028 and beyond. We will provide more details and precise modeling points once we close the transactions.

With that, I will turn to guidance. Our guidance is for stand-alone Palo Alto Networks excluding the pending Applied Compute acquisition. For the fiscal second quarter 2027, we expect NGS ARR between $10.2 billion to $10.22 billion, or 61% growth; RPO between $21.5 billion to $21.6 billion, or 34% to 35% growth; revenue between $3.53 billion to $3.55 billion, or 36% to 37% growth; our fully diluted share count between 837 million to 842 million shares; and our diluted non-GAAP EPS in the range of $1.01 to $1.03.

For the fiscal year 2027, we expect NGS ARR between $11.475 billion to $11.52 billion, or 26% to 27% growth; RPO between $25.3 billion to $25.5 billion, or 19% to 20% growth; revenue between $14.52 billion to $14.57 billion, or 26% to 27% growth; our operating margin of 29.5%; diluted non-GAAP EPS in the range of $4.26 to $4.30; and our fully diluted share count between 844 million to 847 million shares. We expect our adjusted free cash flow margin of 38%.

With that, I will turn it back over to Hamzah for Q&A.

QUESTION AND ANSWER SECTION

Hamza Fodderwala

A

Senior Vice President-Investor Relations, Corporate Development and Strategic Finance, Palo Alto Networks, Inc.

All right. Thank you, Dipak. For our first question, we'll go to Saket Kalia from Barclays followed by Brian Essex from JPMorgan.

Copyright © 2001-2026 FactSet CallStreet, LLC

Saket Kalia

Analyst, Barclays Capital, Inc.

Okay. Great. Hey, guys. Thanks for taking my question here and congrats on a super strong start to the year, really great to see. Nikesh, maybe for you. You've talked about the $1 trillion of tech debt for some time now and maybe that feels even more relevant just given the outages that we've seen. Based on some of your recent customer conversations, where are CEOs and CIOs going to focus that modernization on most in 2027 as the new threat environment gets clearly more active?

Nikesh Arora

Chairman & Chief Executive Officer, Palo Alto Networks, Inc.

Saket, thanks for your question. Look, first and foremost, we're seeing a lot of interest in frontier AI defense. Almost every company is making sure that they are testing their infrastructure with pen testing, red teaming to make sure that all these new AI capabilities, they can inflict them on themselves before bad actors inflict it to them.

I think the challenge, as we highlighted, right now is that there are so many vulnerabilities being discovered on a daily basis that it's impossible for enterprises to patch them. So, I think they're – bit of them are caught in the current sort of crisis at the moment, the short-term. There are people who are beginning to think about how does that impact their long-term architecture, and you can see that. That's reflected in the way we talked about platformizations. We've done 170 this quarter, which means people say, listen, I cannot manage a fragmented infrastructure, how do we get all this stuff put together. The problem is whatever they do today is going to take three to six months, maybe nine months depending on how complex and how tired their environment is. So, they're all embarking on that journey, and I think we're going to see more of that over the next foreseeable future.

Saket Kalia

Analyst, Barclays Capital, Inc.

Very helpful. Congrats again.

Hamza Fodderwala

Senior Vice President-Investor Relations, Corporate Development and Strategic Finance, Palo Alto Networks, Inc.

All right. Thank you, Saket. Next, we'll go to Brian Essex from JPMorgan followed by Rob Owens from Piper Sandler.

Copyright © 2001-2026 FactSet CallStreet, LLC

Brian Essex

Analyst, JPMorgan Securities LLC

Hi. Thank you, Hamza, and thank you for taking the question. Nikesh, for you, would love to – really interesting to see the acquisition. Would love to understand conversations with customers you've had that led to your interest in sovereign AI, and maybe the nature of those conversations and how you're approaching AI sovereignty.

Nikesh Arora

Chairman & Chief Executive Officer, Palo Alto Networks, Inc.

That's a good question, Brian. Look, we started the journey where we said we have built Prisma AIRS. What Prisma AIRS does is it gives you visibility, shows you all the AI artifacts that you have in enterprise. That's easy to do. That's kind of deterministic. You can find it. It's discrete. Then we started to say can we bring all the traffic so we can build an enforcement point and stop all this activity, bad activity from happening. And we figured out the common techniques of LLM sort of hacking through prompt injection, et cetera, and we built controls against that. That was also okay.

When the Hugging Face incident happened, it made us realize that we need to build something other than deterministic controls to intercept AI because deterministic controls cannot anticipate every scenario that an agent is going to conduct, every deception that it's going to try and do. We figured that it was not just a moniker to say fight AI with AI. We actually would have to have an LLM trained to interpret, understand, and stop AI from misbehaving. As we went down that journey, we started looking at, well, how do we go ahead and train models for ourselves. And we went on a journey to look for multiple companies to see who could do it and we looked at the capabilities across the market and Applied Compute stood out.

And then we came to the conclusion that as we train our models for enterprise applications, you will have to train a supervisory AI model at the same time. So, it's almost like having AI build AI and have – build AI that watches this AI. You can't do it independently. I cannot go in and intercept and understand somebody else's application because I don't know how it's supposed to behave because I didn't train it. Imagine if you can train your application and have the same team train something that supervises your application, I think that's the future of AI security.

So, that context, we have to get in the training business to be able to build a good supervision business. That's our bet and we were delighted the Applied Compute team thought that was the right idea. We're excited that we got the deal sort of done and announced today. And I think it's going to be a huge opportunity. In fact, while I'm sitting here, I had a CEO text me saying he met them this morning, and he's really excited now that their concern would have been how do we go to a small team of 100 people and entrust a large enterprise application on them, how they feel with Palo Alto behind them. They can rest assured that this is going to be a long-term play for them. So, I'm very excited. Hopefully, this is the right bet for us, and if it works, I think this will dwarf the abilities that we have in cybersecurity. This could be bigger than what we do today.

Copyright © 2001-2026 FactSet CallStreet, LLC

Brian Essex

Q

Analyst, JPMorgan Securities LLC

Great. I have follow-ups, but I'll leave it for later. Thank you.

Hamza Fodderwala

A

Senior Vice President-Investor Relations, Corporate Development and Strategic Finance, Palo Alto Networks, Inc.

Thank you, Brian. Next, we'll go to Rob Owens from Piper Sandler followed by Meta Marshall from Morgan Stanley. Rob, you're on. Okay. Maybe we'll go to the next person in the queue. We'll go to Meta Marshall from Morgan Stanley followed by Patrick Colville from Scotiabank.

Meta A. Marshall

Q

Analyst, Morgan Stanley & Co. LLC

Great. Thanks for the question. Maybe I just wanted to ask, Nikesh, how you think customers are handling the fact that some of this virtual patching or accelerated patching is clearly causing more outages just as there are errors kind of as they bring some of this online? So, do you see kind of a slowdown? Do you see some form of grouping? Do you see more kind of services involvement? Just how do you see kind of the pickup in outages changing how the patching process is going to go?

Nikesh Arora

A

Chairman & Chief Executive Officer, Palo Alto Networks, Inc.

Going to phone a friend.

Lee Klarich

A

Chief Product and Technology Officer & Director, Palo Alto Networks, Inc.

Thank you, Nikesh. Thank you for the question. Let me first separate the two things you just mentioned. The virtual patching is actually the solution to the challenges associated with patching too quickly and some of the operational impact that comes with that.

Copyright © 2001-2026 FactSet CallStreet, LLC

So, take a step back. All of the vulnerabilities are being found by AI is creating this new wave of patches. That wave of patches then drives the operational challenge of how fast can companies patch everything that they have and that is what is putting the strain on operations teams and that is what then leads to some of the types of outages we've been seeing recently is just organizations trying to move too fast on applying the patches.

What virtual patching allows is for companies to have a mitigation that is safe to deploy, giving their operations teams time to actually apply the real patches. And so, what in essence we're able to provide is the security they need real-time and immediate to the vulnerabilities being discovered, giving them time for the patches to be applied.

And to put this in context, our virtual patching already has thousands of virtual patches as part of the service. By the end of this calendar year, we expect that to be tens of thousands, and it's ramping up very quickly to keep pace with the vulnerabilities that are being discovered.

Q

Meta A. Marshall

Analyst, Morgan Stanley & Co. LLC

Great. Thanks.

A

Hamza Fodderwala

Senior Vice President-Investor Relations, Corporate Development and Strategic Finance, Palo Alto Networks, Inc.

Thank you, Meta. Next, we'll go to Patrick Colville from Scotiabank followed by Shaul Eyal from Cowen.

Q

Patrick Colville

Analyst, Scotiabank

Thanks, Hamza. I guess, Nikesh, and possibly Dipak as well, I want to stick on this theme of the real-world impacts from the heightened cyber risk. In our latest conversations with your largest customers, it's clear that continuous frontier AI defense is like white-hot as a product.

And so, I guess just want to hear a little bit of where you think the scope of this product could be. Is it something that is applicable to all enterprises, mid and large? Is this going to be a category that could mature into like a pretty substantial category or is it a November 2026 kind of pain point seeing spend, but might level out over time? Thank you.

Copyright © 2001-2026 FactSet CallStreet, LLC

A

Nikesh Arora

Chairman & Chief Executive Officer, Palo Alto Networks, Inc.

Hey. There has been a notion of vulnerability management historically, which was an industry TAM, which was more sort of a set of professionals finding vulnerabilities, institutionalizing them and selling them as feeds for others to be able to go look at the infrastructure and match them.

Now, what's happened is with the capabilities we've seen with frontier LLMs in the last few months is that it's become an institutionalized capability, which you don't need to go send outsource, and you can custom check your infrastructure and configuration. Unfortunately, it's in such a way that even bad actors can do it. So, there's a mad rush for every enterprise to go make sure that they have a robust perimeter, robust infrastructure to make sure they don't get breached. There's a mad rush for vendors to make sure that they can plug any holes from a vulnerability perspective.

I think the progression the last six months has been such that a lot more are getting found now with the most recent models than there were six months ago. So, I don't think the capabilities of model has plateaued yet and I don't expect it will plateau for another six to nine months, which means you will keep getting better models over the next six to nine months, smarter models which will find much more – sort of can do much more from a simulation perspective, find your [ph] attack pots (00:36:07). So, I think for the next 12 months or pick your number, 12 to 18 months, there is going to be the need for persistent frontier AI testing, which is why we call it continuous frontier AI defense because we don't think it's going to stop in the short-term.

I do believe from an industry structure perspective over time, the continuous frontier AI defense will morph into a capability that every company would have delivered to them by their platform providers, whether it's through our XSIAM platform, through our perimeter security platform because our platforms will have to constantly check and stress test their infrastructure and start self-healing over time as opposed to expecting them to find them and then get – talk to us about fixing it.

So, you'll feel all of us in the industry will absorb that into our capabilities, but I think that's two years out. And until then, there is going to be this constant desire or need for customers to use either open-weight, open-source models or frontier LLMs to test their infrastructure.

Q

Patrick Colville

Analyst, Scotiabank

[ph] Okay (00:37:02).

Copyright © 2001-2026 FactSet CallStreet, LLC

Hamza Fodderwala

A

Senior Vice President-Investor Relations, Corporate Development and Strategic Finance, Palo Alto Networks, Inc.

Thank you, Patrick. Next, we have Shaul Eyal from Cowen followed by Fatima Boolani from Citi. Shaul?

Patrick Colville

Q

Analyst, Scotiabank

Yeah. Hey, good afternoon. Congrats, guys. So, firewall bookings grew 28% on a TTM basis. How much reflects AI-related demand versus refresh cycle? What supports continued growth? So, maybe in other words, Nikesh, the cyclical versus the structural question, cyclical versus structural demand?

Nikesh Arora

A

Chairman & Chief Executive Officer, Palo Alto Networks, Inc.

First and foremost, you have to understand hardware is a very small part of our business, right? And our product number includes both software capability as well as hardware. So, it's not pure hardware that you're seeing in there and the software firewalls are being driven by AI security. Just keep that in mind as you think about the balance between product of hardware and software.

Having said that, yes, we have seen an uptick in hardware as the industry has seen as well. I think part of it is people accelerating purchases because prices are going up because of component pricing. Component pricing is high across the industry, as you know, because of the AI sort of factor. There is increase in demand. There is a bunch of neoclouds which are needing to make their infrastructure robust over time. So, you're seeing that we are sort of benefiting from the neocloud build cycle and I think there is an element of consolidation. People want to start ripping out part of the old infrastructure and replacing it. So, I think we're seeing all three effects because we've never seen this level of hardware growth historically. We think this persists at least the next 12 months, if not longer.

Patrick Colville

Q

Analyst, Scotiabank

Thank you so much.

Copyright © 2001-2026 FactSet CallStreet, LLC

Dipak Golechha

A

Chief Financial Officer, Palo Alto Networks, Inc.

If I can maybe just add one other piece is what we're also seeing is we're seeing customers that may have said, look, let me have two vendors like just say, look, it's much easier for me to have one vendor so that I can get all the patching and everything else done better as well. So, multiple different factors, but I think the overall point is we think this is very, very robust and will continue.

Shaul Eyal

Q

Analyst, TD Cowen

Thank you.

Hamza Fodderwala

A

Senior Vice President-Investor Relations, Corporate Development and Strategic Finance, Palo Alto Networks, Inc.

All right. Thank you, Shaul. Next, we'll go to Fatima Boolani from Citi followed by Jonathan Ho from William Blair.

Fatima Boolani

Q

Analyst, Citigroup Global Markets, Inc.

Thank you. Good afternoon. I appreciate you taking my question. Nikesh, I wanted to go back to the impetus and the rationale for Applied Compute. I think it's pretty clear what the strategy is, certainly to me, but I wanted to get a sense from you on how this evolves or diminishes or changes your value proposition as a Switzerland-esque cybersecurity platform to the – an AI infrastructure platform ecosystem because necessarily this could develop into a competitive co-optative relationship with some of your customers and partners.

And then as a related matter, maybe for Lee, to what extent should we think about the outputs of Applied Compute being fed back into the business for the purposes of generating new SKUs and new capabilities and products that burnish your own platform and your security offerings and capabilities.

Nikesh Arora

A

Chairman & Chief Executive Officer, Palo Alto Networks, Inc.

Copyright © 2001-2026 FactSet CallStreet, LLC

I'm trying to parse Switzerland and co-optative and what's that, competitive. All right, good. Let me parse that. Okay.

Q

Fatima Boolani

Analyst, Citigroup Global Markets, Inc.

That's it.

A

Nikesh Arora

Chairman & Chief Executive Officer, Palo Alto Networks, Inc.

First and foremost, Fatima, look, there is not enough AI skill in the market at frontier lab level that any one of us have. So, just at the outset, the basic thing these people will do is they will significantly increase our AI capabilities as an organization because these people have been working at frontier labs and that's the mindset they're arriving in. We've been working hard to train our products with LLMs and I think this team will supercharge and turbocharge our abilities, allowing us to deliver AI capabilities to our customers. So, that's number one.

If we do nothing else, that's a win. I think the second part, as we said, we discovered that there is a market out there where we looked for this. We looked at nine different providers and we figured out these guys are the best at training, but even the best of them were still in their early stages evolving as enterprises and organizations.

And we know from our life in enterprise that customers want a reliable partner who has all the bells and whistles to be able to provide accountability, to provide capability at speed. So, I think enveloping them with our capabilities from a go-to-market perspective, resilience perspective, cybersecurity perspective, I think, is going to expand their existing business itself without doing anything else. Again, that's a win.

I think the third part, which excites me the most is as we go down the journey of AI security, I'm coming to the conclusion that before you start building a Switzerland product, let's make sure your Apple product works well, right? Because right now, the end-to-end capability does not exist in the industry. You need to make sure you have the same vendor doing discovery of your AI assets, the same vendor aggregating that traffic into a choke point, the same vendor understanding the activities of an LLM so we can decide what to allow, what to block.

I think the idea that I can deliver a security product to somebody who's built their own AI intelligence, and I have no visibility into what's happening in there yet I can somehow interpret what that AI is trying to do and try and stop it is insane. It's not going to happen and as you can see, there's a whole science developing in AI around evals, around traces. If you saw most recently, two days ago, OpenAI announced that they have done the deal with four different companies who are going to have early access to their models so that they can understand how the models are behaving and if the models misbehave, they'll be able to tell OpenAI how to train them better.

Copyright © 2001-2026 FactSet CallStreet, LLC

That's the level of intelligence we need to be able to stop AI. If I don't understand that the model is behaving or not behaving, how am I going to be able to stop it? For that, I need frontier LLM capability, frontier LLM skills. That's what makes sense. That's why we're doing what we're doing. I think that's what you will see that eventually, that's where the industry's going to want to go. And you'll see a lot of people try and follow our lead on this one. And then we'll [ph] build LLM (00:43:17) to it don't worry.

Q

Fatima Boolani

Analyst, Citigroup Global Markets, Inc.

I appreciate it. Thank you.

A

Hamza Fodderwala

Senior Vice President-Investor Relations, Corporate Development and Strategic Finance, Palo Alto Networks, Inc.

Thank you, Fatima. Next, we'll go to Jonathan Ho from William Blair followed by Michael Turrin from Wells Fargo.

Q

Jonathan Ho

Analyst, William Blair & Co. LLC

Excellent. Good afternoon. Congratulations on the strong results. With Prisma AIRS, this is now your fastest product to $170 million. Is there a way to understand where we are in the adoption curve of AI and do you see 2027 as maybe a potential inflection point for this to accelerate even further as folks move from developer use cases to now true enterprise agentic and inference use cases? And maybe help us understand why the demand curve changes as you change between these new use cases. Thank you.

A

Nikesh Arora

Chairman & Chief Executive Officer, Palo Alto Networks, Inc.

[indiscernible] (00:44:11). I think it's very early. As I said, if you ask any enterprise, [ph] you got to do a lot of it. You ask them (00:44:15), are you using coding applications, they'll tell you, they are. Ask them are they secure and the answer's going to be, what do you mean? People haven't figured out how to secure Claude Code or Codex or Cursor yet. I think that's the starting point.

Copyright © 2001-2026 FactSet CallStreet, LLC

You got to make sure as these coding applications get agentic capabilities, they're going to be [ph] respond (00:44:33) agents. How are you going to control them? And I don't think most enterprises have figured that problem out and they're busy using coding tools.

And I expect that you will see some of those agents created as part of coding tools misbehave. I think that's where we start. That's the first application of Prisma AIRS is don't let your [indiscernible] (00:44:50) coding agents run amok in your enterprise, but I'm going to let Lee talk about future use cases that applies.

Lee Klarich A

Chief Product and Technology Officer & Director, Palo Alto Networks, Inc.

Yeah. I think first of all, I agree with Nikesh that this is still very...

Nikesh Arora A

Chairman & Chief Executive Officer, Palo Alto Networks, Inc.

Good strategy.

Lee Klarich A

Chief Product and Technology Officer & Director, Palo Alto Networks, Inc.

We're still very early. Companies are still very early in their AI adoption. The AI coding use case is really the first one that has gotten across the line to being somewhat common from a production perspective. The vast majority of other AI and agentic use cases are still very much in the experimentation, piloting, development stage.

And what's often blocking them from moving to full production is, number one, can they be secured; and two, do they perform accurately. And I think both of those challenges can be solved. And I think part of what we're doing with Applied Compute will actually provide a significant part of the answer to how we can get solved both of those problems to address the broader enterprise use cases of adopting agents and AI overall.

Jonathan Ho Q

Analyst, William Blair & Co. LLC

Thank you.

Copyright © 2001-2026 FactSet CallStreet, LLC

Hamza Fodderwala

A

Senior Vice President-Investor Relations, Corporate Development and Strategic Finance, Palo Alto Networks, Inc.

Thank you for the question, Jonathan. Next, we'll go to Michael Turrin from Wells Fargo followed by Josh Tilton from Wolfe Research.

Michael Turrin

Q

Analyst, Wells Fargo Securities LLC

Thank you, Hamza. Nikesh, it looks like some of the warning signs you were flashing are now here. So, maybe speak to where the budget comes from an enterprise, how much consolidation of the platformization strategy helps drive those discussions forward, and if the company is tooled to handle a persistent surge in demand or how yourself and the industry just work through all of this going forward. Thank you.

Nikesh Arora

A

Chairman & Chief Executive Officer, Palo Alto Networks, Inc.

I think the frontier AI defense budgets have just come out of emergency funds. I think there is no leadership team that has sat there and says we don't want to do this. We can't do this. They said we have to do this. So, I don't know where they found them, but they actually found the money to make sure they can check for vulnerabilities.

And I think as you see, just the events the last two weeks, if somebody has been sleeping under a rock, they must have come out and said, holy [obscenity] (00:46:53). If we are not doing this, we better go out and test ourselves. So, I'm not worried about the fact that we're not. We're going to see people show up with money.

I think the longer term question is still live. A lot of CIOs are – CEOs are going to see us saying, listen, how do we fix this for the long-term and I think that's why you're seeing consolidation. That's where you're seeing sort of how do I make – sweat my dollar further. How do I make sure that I'm not spending inconsistently with where I want to be in the long-term?

So, you'll see cyber budgets rise, I'm sure, in the next two years. You will also see a degree of consolidation because the spray and pray strategy has not worked in cyber.

Michael Turrin

Q

Analyst, Wells Fargo Securities LLC

Thank you.

Copyright © 2001-2026 FactSet CallStreet, LLC

Hamza Fodderwala

A

Senior Vice President-Investor Relations, Corporate Development and Strategic Finance, Palo Alto Networks, Inc.

All right. Thank you, Michael. Next, we'll go to Josh Tilton from Wolfe Research followed by Gabriela Borges from Goldman Sachs.

Joshua Tilton

Q

Analyst, Wolfe Research LLC

Hey, guys. Can you hear me?

Hamza Fodderwala

A

Senior Vice President-Investor Relations, Corporate Development and Strategic Finance, Palo Alto Networks, Inc.

Yes. We can hear you, Josh.

Joshua Tilton

Q

Analyst, Wolfe Research LLC

Awesome. I kind of want to touch on a comment I heard in the prepared remarks. I think you guys mentioned that agentic traffic on SASE took a decent spike in the quarter or over some recent time period. I think when we talk to investors, I don't really think they view SASE as kind of this market that maybe what you would call an AI winner.

Can you guys talk to what you're seeing in kind of this agentic opportunity, specifically in SASE? Clearly, you have something going on the platform today and just maybe one level deeper, are we monetizing these agents on the SASE platform? Are they incremental to the human traffic? How do we think about that opportunity specifically for SASE?

Nikesh Arora

A

Chairman & Chief Executive Officer, Palo Alto Networks, Inc.

Copyright © 2001-2026 FactSet CallStreet, LLC

I thought you wanted to understand industry, but you really want to know money. Okay. So, look, historically SASE has been a per user pricing model, right? Because the traffic per user has a capacity limit, human beings can only transact so much traffic over a period of a month or a period of years.

So, SASE was sold on a per user, per seat basis. You saw – you can see as an industry parallel that traffic to WAFs has increased because there's a bunch of agents running around in the public domain trying to go drive traffic.

So, imagine the same activity happening inside the enterprise. You have agents that are going to start running around enterprises as well, and that's going to put load on the SASE fabric.

To embrace that, I think we have done this and I suspect other industry are also doing that. We used to sell a certain amount of free network capacity as part of SASE, but we're going to get more diligent about metering it because at the end, it increases COGS if you drive more traffic through our SASE fabric.

So, you will see us monetize the fabric. The traffic, right now we've seen a 9x spike. Even at 9x, it's sort of on an organizational level, it's not increased compared to [ph] what our (00:49:30) licensing figures, but I expect in a short period of time, they will start blowing past the total capacity of licensed capacity of traffic in SASE. And as they see more agents in the enterprise, we will see SASE will start seeing some traffic monetization, but I suspect it's still six to nine months out. But it's going to happen.

A

Hamza Fodderwala

Senior Vice President-Investor Relations, Corporate Development and Strategic Finance, Palo Alto Networks, Inc.

All right. Thank you, Josh. Next, we'll go to Gabriela Borges from Goldman Sachs followed by John DiFucci from Guggenheim.

Q

Gabriela Borges

Analyst, Goldman Sachs & Co. LLC

Hey. Good afternoon. Thank you. Nikesh, I want to follow-up on a comment you've made in the recent past on wanting to be more in the token flow and certainly Chronosphere and observability fits that.

Certainly, I would think something like Applied Compute fits that. Maybe just zoom out for us. As you see more AI adoption, do you think more categories are going to intersect with security? And talk to us a little bit about what that convergence of security and infrastructure looks like.

Copyright © 2001-2026 FactSet CallStreet, LLC

A

Nikesh Arora

Chairman & Chief Executive Officer, Palo Alto Networks, Inc.

Look, if you look at our historical business model for cybersecurity, it's been more of an enterprise license model. It's also been a per user base model. As you know, per user base models are going to morph because of agents that come into play.

And the enterprise license model had, in certain cases, a sort of a volume orientation. So, for example, software firewall is volume-oriented. The more traffic you send through them, the more capacity is needed, the more you have to buy.

Prisma AIRS is exhibit A. Prisma AIRS is fundamentally based on inspecting tokens. The more tokens you use in the enterprise, the more you have to have inspection tokens, inspection tokens cost, and Prisma AIRS is possibly going to be 2% to 5% of your spend over the long-term of your token spend.

If you believe there's $1 trillion of tokens that's going to be sold in the next few years, then by definition, the industry must see a 3% to 5% security spend associated with tokens. I think the same thing applies to Applied Compute, no pun intended. You've seen that you yourself said it, that Applied Compute is effectively in the training business. The more models you train, the more models you run through inference, but on top of that, Applied Compute is going to come with a security capability in the future where you'll have to pay both for training your regular application LLM, and you'll have to pay for security.

I think in any one of our products where we're going to see increased traffic, increased tokens, we're going to get in the token flow. So, if you look at our three businesses from a network security, from a Cortex and from identity and observability and the new business of Applied Compute, observability and Applied Compute are both trafficdetermined. The more traffic that flows to the AI ecosystem, observability wins. The more traffic that is needed in the AI ecosystem, Applied Compute wins.

The more traffic that is run through an organization, Cortex wins because you ingest it. So, the more traffic you see, Prisma AIRS and network security wins because you need bigger firewalls and software, hardware, and SASE. So, the only place where you don't see a direct correlation outside of enterprise licensing is Idira, but Idira also that's why we're building agentic capabilities. The more number of agents you see is going to drive more non-human identity in Idira. So, we're pivoting to ensure all of our businesses are aligned towards token flow because over time, as AI gets more and more ubiquitous and as it gets more and more consumptive, we think Palo Alto's business model will align with the consumption of AI.

Q

Gabriela Borges

Analyst, Goldman Sachs & Co. LLC

I like the framing. Thank you.

Copyright © 2001-2026 FactSet CallStreet, LLC

Hamza Fodderwala

A

Senior Vice President-Investor Relations, Corporate Development and Strategic Finance, Palo Alto Networks, Inc.

[ph] Okay (00:52:52). Thank you, Gabriela. And we'll take our last question from John DiFucci at Guggenheim.

John DiFucci

Q

Analyst, Guggenheim Securities LLC

Thanks, Hamza, and nice job, everybody. This is really strong results. Everybody can see that, but this acquisition is so different from what you do today in a lot of ways and that became apparent with Dipak's preparatory comments.

But Nikesh, is the compute realm something we should think about you personally having experience with given your experience prior to Palo? And Dipak, how should we think about the financial profile of Palo Alto over time, especially if this business matches your aspirations?

Nikesh Arora

A

Chairman & Chief Executive Officer, Palo Alto Networks, Inc.

[ph] John (00:53:31), as you know, I made a career out of doing things which I don't have experience in.

John DiFucci

Q

Analyst, Guggenheim Securities LLC

Okay.

Nikesh Arora

A

Chairman & Chief Executive Officer, Palo Alto Networks, Inc.

So, you can rely on my lack of experience in compute to make me win in compute. Sorry, that was a joke because he asked about my experience in compute. It's so hard to understand compute.

Copyright © 2001-2026 FactSet CallStreet, LLC

You take a large piece of land, you make sure it's got power, you make sure that it's got water so you can air cool the – water-cool the chips. You make sure you get some chips from one of your friends and you go out and build a large data center. Well, good news is there's a lot of people doing that building right now.

So, Applied Compute currently has a bunch of compute contracts available with these large players. In addition to that, obviously we are a large consumer of AWS compute and large consumer of GCP compute. We have not had to have compute experience to run our capabilities so far. Maybe we don't and maybe we will do some inference partnerships to get more compute as we grow.

[indiscernible] (00:54:21). The first thing in my mind in Applied Compute is generate demand. Generating demand requires us to convince enterprises that they need to go with Palo Alto and apply compute because they need to get secure inference delivered to them and secure training. Once we – and we can build that on pay-as-you-go compute. We can build that on three-year contracts in compute, which are already in [ph] places like (00:54:41) compute doesn't require us to become compute experts.

Of course, as we see demand grow, if you want to extract more margin from the compute line, then we can make sure that we get into the risk business of aggregating compute. That becomes a different conversation. Again, we've said that we aspire to have 100 to 150 megawatts of compute.

I can do that with a neocloud. I can do that with cloud providers out there. I don't need to be an expert yet. If you go past that and you start talking gigawatts, yes. Then I have to worry about being a compute specialist. Until then, I think we can manage. I think the key here is to drive more top line growth, which is going to be driven by secure delivery of inference to the enterprise and training.

John DiFucci

Analyst, Guggenheim Securities LLC

Q

Okay. Thanks. And Dipak, how's that going to affect the financial profile over time?

Dipak Golechha

Chief Financial Officer, Palo Alto Networks, Inc.

A

Look, I think the most important part, John, is to understand that we're taking a capital-light approach, right? So, I think we do have many options. I think I talked about in my prepared remarks what we think the impact is going to be on our financials in the short-term and the medium term and more broadly in the long-term.

But I think our North Star is to make sure we provide the security profile that's needed for enterprises to act, and that's where we think the money will be made. And that continues to be a profit zone. Then we'll just have another segment that's very profitable.

Copyright © 2001-2026 FactSet CallStreet, LLC

Q

John DiFucci

Analyst, Guggenheim Securities LLC

Well, you guys have done a lot of – made a lot of bold moves and so far, at least from my observation, they've been pretty right on. So, we'll see. We'll see how this goes.

A

Dipak Golechha

Chief Financial Officer, Palo Alto Networks, Inc.

Thank you.

Hamza Fodderwala

Senior Vice President-Investor Relations, Corporate Development and Strategic Finance, Palo Alto Networks, Inc.

Thank you, John. That concludes the Q&A portion of the call. I'll turn it back to Nikesh for concluding remarks.

Nikesh Arora

Chairman & Chief Executive Officer, Palo Alto Networks, Inc.

I just want to say thank you to all of our employees and our partners and our customers for trusting in us and employees for doing a lot of hard work to get us here. This is a great start to our fiscal year, and I hope we continue to work hard to deliver again. Thanks, everyone. Thanks for joining the call.

Disclaimer

The information herein is based on sources we believe to be reliable but is not guaranteed by us and does not purport to be a complete or error-free statement or summary of the available data.

As such, we do not warrant, endorse or guarantee the completeness, accuracy, integrity, or timeliness of the information. You must evaluate, and bear all risks associated with, the use of any information provided hereunder, including any reliance on the accuracy, completeness, safety or usefulness of such information. This information is not intended to be used as the primary basis of investment decisions. It should not be construed as advice designed to meet the particular investment

Copyright © 2001-2026 FactSet CallStreet, LLC

needs of any investor. This report is published solely for information purposes, and is not to be construed as financial or other advice or as an offer to sell or the solicitation of an offer to buy any security in any state where such an offer or solicitation would be illegal. Any information expressed herein on this date is subject to change without notice. Any opinions or assertions contained in this information do not represent the opinions or beliefs of FactSet CallStreet, LLC. FactSet CallStreet, LLC, or one or more of its employees, including the writer of this report, may have a position in any of the securities discussed herein.

THE INFORMATION PROVIDED TO YOU HEREUNDER IS PROVIDED "AS IS," AND TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, FactSet CallStreet, LLC AND ITS LICENSORS, BUSINESS ASSOCIATES AND SUPPLIERS DISCLAIM ALL WARRANTIES WITH RESPECT TO THE SAME, EXPRESS, IMPLIED AND STATUTORY, INCLUDING WITHOUT LIMITATION ANY IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, ACCURACY, COMPLETENESS, AND NON-INFRINGEMENT. TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE
LAW, NEITHER FACTSET CALLSTREET, LLC NOR ITS OFFICERS, MEMBERS, DIRECTORS, PARTNERS, AFFILIATES, BUSINESS ASSOCIATES, LICENSORS OR SUPPLIERS WILL BE LIABLE FOR ANY
INDIRECT, INCIDENTAL, SPECIAL, CONSEQUENTIAL OR PUNITIVE DAMAGES, INCLUDING WITHOUT LIMITATION DAMAGES FOR LOST PROFITS OR REVENUES, GOODWILL, WORK STOPPAGE, SECURITY BREACHES, VIRUSES, COMPUTER FAILURE OR MALFUNCTION, USE, DATA OR OTHER INTANGIBLE LOSSES OR COMMERCIAL DAMAGES, EVEN IF ANY OF SUCH PARTIES IS ADVISED OF THE POSSIBILITY OF SUCH LOSSES, ARISING UNDER OR IN CONNECTION WITH THE INFORMATION PROVIDED HEREIN OR ANY OTHER SUBJECT MATTER HEREOF.

The contents and appearance of this report are Copyrighted FactSet CallStreet, LLC 2026 CallStreet and FactSet CallStreet, LLC are trademarks and service marks of FactSet CallStreet, LLC. All other trademarks mentioned are trademarks of their respective companies. All rights reserved.