01-Sep-2026

Palo Alto Networks, Inc. (PANW)

Q4 2026 Earnings Call

CORPORATE PARTICIPANTS

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

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

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

OTHER PARTICIPANTS

Rob D. Owens
Analyst, Piper Sandler & Co.

Brian Essex
Analyst, JPMorgan Securities LLC

Saket Kalia
Analyst, Barclays Capital, Inc.

Fatima Boolani
Analyst, Citigroup Global Markets, Inc.

Matthew Hedberg
Analyst, RBC Capital Markets LLC

Michael Turrin
Analyst, Wells Fargo Securities LLC

Gray Powell
Analyst, BTIG LLC

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

Brad Zelnick
Analyst, Deutsche Bank Securities, Inc.

MANAGEMENT DISCUSSION SECTION

Hamza Fodderwala

Senior Vice President-Investor Relations & Strategic Finance, Palo Alto Networks, Inc.

Good day, everyone, and welcome to Palo Alto Networks Fiscal Fourth Quarter 2026 Earnings Conference Call. I am Hamza Fodderwala, Senior Vice President of Investor Relations and Strategic Finance. Please note that this call is being recorded today, Tuesday, September 1, 2026 at 1:30 PM Pacific Time.

With me on today's call to discuss our fiscal fourth quarter results are Nikesh Arora, our Chairman and Chief Executive Officer; and Dipak Golechha, our Chief Financial Officer. 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 Q4 2026 Supplemental Financial Information and Q4 2026 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.

I will now turn the call over to Nikesh.

Nikesh Arora

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

Thank you, Hamza. Good day, everyone, and thank you for being with us to discuss our progress. As you can see, our execution fueled record finish to the fiscal year. We exceeded our guidance across every financial metric in Q4, with bookings momentum accelerating for the second straight quarter. This performance is a direct result of record-breaking platformization adoption and the growing urgency among customers to fortify their defenses as AI fundamentally redefines the security landscape.

We achieved record RPO, surpassing the $20 billion threshold for the first time, to close the year at $21.2 billion, representing a growth rate of 34%. NGS ARR reached $9.1 billion, up 63%, enabling us to report one of our most substantial Next-Generation Security ARR outperformances to-date. Most notably, we added nearly $1 billion in net new NGS ARR this quarter alone.

I remember my first Analyst Day in 2019, shortly after I arrived, we set a high bar to reach $1 billion in Next-Generation Security revenue by fiscal 2022, just as we were initiating our pivot from a single-product firewall

vendor into a unified security platform. That transformation journey has reached a pivotal inflection point, and the scale of our current success is a testament to that vision.

We delivered broad-based strength across our platforms in Q4, with Network Security, our largest business, reporting exceptional results across SASE, software, and hardware firewalls. XSIAM maintained its strong momentum, while Prisma AIRS achieved a significant milestone, surpassing $100 million in ARR within four quarters of general availability. This represents the fastest scaling product in the history of Palo Alto Networks.

Fiscal 2026 marked a pivotal inflection point in our transformation journey. We closed the two largest acquisitions in our history with CyberArk and Chronosphere, both of which are exceeding our initial expectation. Both businesses are gaining significant traction within our platformized architecture and are scaling at an accelerated pace compared to their previous stand-alone performance. These achievements are a testament to the execution and deep collaboration of the thousands of new colleagues who joined us this past year. We look forward to continuing this shared momentum into FY 2027.

Q4 was the very first quarter in which we witnessed the profound implications of cyber-capable models. As I've said before, AI is a long-term tailwind for cybersecurity. While these models are becoming increasingly proficient at uncovering vulnerabilities, detection is merely the opening act. Truly validating, interpreting context, and resolving these issues requires broad cybersecurity platforms working alongside frontier AI. This synergy is essential to stress test environments, manage agentic actions, and trigger machine-speed remediation during an active threat.

Defending at that speed necessitates a unified data architecture where AI processes every signal, collapsing response times from days to just minutes. Platformization is the only viable strategy for real-time defense, fighting AI with AI, and that philosophy continued to gain significant resonance with our customers in Q4.

During the fourth quarter, we achieved approximately 220 net new platformizations surpassing our prior record and representing more than twice the volume for when we initiated this metric two years ago. The performance validates that our philosophy of real-time defense through unified architecture continues to gain significant resonance.

Beyond initial adoption, standardizing on our platform yields superior retention and expansion with NRR, or net revenue retention, exceeding 120% for our platformized cohort in Q4. As we look forward, we remain on track towards our long-term objective of over 4,000 platformizations by fiscal 2030, which serves as a bedrock for reaching our $20 billion Next-Generation Security ARR target.

Our largest Q4 wins show platformization in action. During the fourth quarter, we secured a $126 million agreement with a global telecoms leader. This organization moved to standardize on our Network Security platforms, bolstering their next-generation firewall footprint while displacing legacy proxy providers with Prisma Access for SASE.

We also closed a $72 million transaction with a premier IT service provider. This client has fully embraced platformization across Network Security, Cortex, and Idira, making eight-figure investments in each, serving as a powerful validation of our cross-sell momentum in Q4.

A further highlight was a $53 million platformization deal with a leading global payments platform. Beyond standardizing their network defense on our architecture, they committed high-seven-figures to Prisma AIRS as they accelerated their enterprise AI initiatives.

Fiscal 2026 has emerged as a landmark period in the rapid evolution of AI, marked by three distinct inflections over the last six months. Each of these shifts fundamentally redefines how AI interacts with enterprise and, by extension, how it impacts the cybersecurity landscape. For us to effectively lead and protect our customers, maintaining our position as a vanguard of these structural changes is paramount.

The first inflection was the arrival of OpenClaw. Earlier this year, OpenClaw served as the catalyst for the transition from standard LLMs to agentic action, fundamentally altering the dynamic between human operators and AI systems. Just a year ago, AI was largely defined by individual human prompting, a synchronous, multi-turn dialogue, where task was completed with a person in the loop. Virtually overnight, we witnessed the emergence of fully autonomous agents. These are persistent entities that operate for extended durations, executing complex workflows without direct supervision. For a single employee who once managed one task at a time, that same individual can now orchestrate thousands of autonomous agents.

The implications for the enterprise are profound. Each of these agents generates continuous traffic, interacting with models, creating internal data, and communicating with other tools and agents around the clock. This creates a massive volume of telemetry that must be observed, while every agent requires its own set of credentials. We're now securing a whole new class of machine identities with autonomous permissions. The surge in traffic, data, and identity complexity represents a significant long-term tailwind across every one of our platforms.

The second was the Mythos moment, which proved that deep domain training enables AI to achieve unprecedented proficiency. In our sector, this has manifested as the weaponization of AI to identify and exploit vulnerabilities at scale. This shift has exposed the deep technical debt within the enterprise, where legacy flaws and persistent misconfigurations that once took months for a human to uncover are now exploited in minutes.

In an AI-driven threat environment, there is no longer anywhere to hide. For our customers, the Mythos moment reframed the security challenge from visibility to velocity. Organizations must now identify exposures before they are weaponized and respond at machine speed. This is why real-time defense has shifted from a future road map item to a present-day requirement.

To address this, we expanded our Frontier AI Defense Service last month, introducing a multimodal harness that enables enterprises to stress-test their environments. This service leverages the most sophisticated cyber-capable models available, and we are proud to be the first certified commercial partner for Mythos 5.

The third involves an emerging inflection point that we expect will dominate the cybersecurity dialogue in the coming quarters. For the past 90 days, the market has moved beyond a handful of frontier models towards a diversified ecosystem of open-weight and open-source architectures. Enterprises are increasingly prioritizing sovereign control over their AI, leading to the deployment of specialized models deeply integrated with proprietary data.

We expect a major acceleration as organizations utilize internal telemetry to fine-tune models for bespoke enterprise use cases. While frontier models will continue to set the high watermark for intelligence, the broader market is heading towards rapid fragmentation and proliferation. Crucially, each new deployment adds more infrastructure to fortify and more sensitive data to protect. The surface area requiring platformized protection is expanding dramatically.

Three pivotal moments, each with a unique impact, yet all leading to a single conclusion. As the relationship between humans and AI evolves and deployments multiply, the necessity for unified, real-time defense has never

been greater. It is early days, but we are beginning to see the signs of how these trends are impacting our business.

Starting with our largest business, Network Security. AI represents a significant long-term tailwind that is expanding our total addressable market in network security while reinforcing that platformization is the only viable strategy for the modern enterprise. As the global AI buildout continues, every new data center becomes critical infrastructure that requires robust fortification through hardware and software firewalls, whether delivered natively by cloud providers or via a unified security platform.

The ecosystem driving this infrastructure expansion has reached a pivotal inflection point, and now we're seeing a new vanguard of buyers emerge, spanning sovereigns, neoclouds, and frontier AI labs, all racing to deploy massive computational capacity that must be secured. We achieved strong early traction with this cohort in FY 2026, including multiple seven-figure bookings in the fourth quarter.

In total, our firewall execution drove accelerated bookings through the fiscal year, fueled by robust demand for our latest Gen 5 hardware and the continued momentum of our software offerings as customers scale their cloud and AI workloads. As this infrastructure matures and autonomous agents are deployed, we expect a dramatic proliferation of agentic traffic across every network and cloud environment.

The impact on our SASE platform is already evident, where agentic traffic has surged 9x over the last nine months. Defending at this scale requires machine-speed inspection, a core competence we have refined for two decades, enabling us to block more than 30 billion attacks in a single day. Ultimately, AI is underscoring the urgent need for unified platforms that deliver real-time defense.

In FY 2026, our platform advantage drove exceptional results in our SASE businesses, where bookings grew 40% with broad strength across access, SD-WAN, and secure browser. We successfully displaced legacy incumbents in nearly 100 accounts, representing over $400 million in total contract value, nearly double the volume of displacement from a year ago.

While we have rapidly ascended to the number two position in this market, we're playing to win and remain on a clear trajectory to become the SASE leader within the next five to seven years. We're in the early chapters of this shift where the future necessitates securing both human and machine identities through unified architecture capable of providing defenses at machine speed.

Organizations are transitioning AI initiatives from experimentation to full-scale production, significantly widening the defensive perimeter with each new deployment. Prisma AIRS has continuously adapted alongside these adoption cycles, evolving to mitigate the unique risks emerging from every phase of the AI journey.

While our initial focus addressed the chatbot-centric era of generative AI, our vision has expanded towards providing a comprehensive architecture for agentic security. This unified approach begins with securing machine identities and credentials, incorporates deep observability of agentic footprints, and extends to the endpoint, where we analyze behavioral intent. By funneling this traffic through our AI gateway, we ensure that security policies are enforced in real time across every interaction.

Prisma AIRS achieved a remarkable milestone in Q4, surpassing $100 million in ARR within just four quarters of general availability, marking the most rapid scale out of any product in our history. Our momentum is reflected in a growing base of over 800 customers for this product, with the majority of our largest transactions now featuring multi-module adoption in Q4.

We're also seeing significant early validation of our agentic endpoint strategy following the Koi acquisition. We believe the endpoint is reaching a critical inflection point as AI development tools migrate to the desktop environment. This shift creates an expanded surface area where agents autonomously manage files and access sensitive credentials.

Legacy security tools often remain blind to the underlying intent and reasoning behind these machine-speed actions. In this landscape, visibility without action is insufficient. Our platformized approach delivers end-to-end transparency from the initial prompt to the final execution, enabling in line prevention at machine speed.

This capability is becoming a fundamental requirement for the enterprise. We've already secured over 100 logos, representing a 2.5 times increase since finalizing the Koi integration earlier this year. Ultimately, the synergy of detection and prevention is most effective when unified as a single platform, with XSIAM serving as a central nervous system for this critical telemetry.

Earlier this year, our Unit 42 researchers demonstrated the staggering speed of modern threats by simulating a comprehensive AI-driven attack in under 30 minutes. Contrast that with the industry standard defense report response of four days, and it's clear that legacy approaches are no longer sustainable.

Customers standardizing on XSIAM are transforming their operation, reducing their mean time to respond to less than 10 minutes, massive leap from the days or weeks required previously, as we continue our relentless push towards true real-time defense. In the fourth quarter, XSIAM maintained its exceptional momentum, concluding the year with over $700 million in ARR, up 70%, while surpassing the 1,000th customer milestone on the platform.

The power of our architecture lies in the fact that live telemetry is already resident within XSIAM, allowing us to seamlessly unlock new value through our unified data lake. Expanding a deployment does not require the friction of new product integration. It simply involves querying existing data in new ways. As of Q4, the majority of customers have embraced this platform advantage, utilizing multiple modules, including exposure management and cloud security.

Turning to Observability. We continue to see the world's premier AI-native and cloud-first organizations standardize on our technology. The entire – entities pioneering the AI frontier generate telemetry at a scale that traditional tools cannot withstand. Chronosphere has engineered specifically for these massive data volumes, capturing every training run and agent loop.

This quarter, we signed a $20 million deal with a hyper-growth AI inference provider that processes tens of trillions of tokens a day. There is no stronger validation of our platform than when the architects of the AI ecosystem trust us to monitor their own infrastructure. Since finalizing the Chronosphere acquisition in Q2, our Observability ARR has more than doubled, eclipsing the $500 million mark. This performance has significantly outperformed our initial targets and represents the most rapid post-acquisition scaling in our history. Our cross-sell strategy is delivering tangible results, with XSIAM contributing to 50% of net new Chronosphere logos this quarter through multiple seven-figure agreements.

We are further enriching the stack with the acquisition of Embrace, integrating real user monitoring to complement our core metrics, logs, and traces. This expansion enables us to provide a comprehensive end-to-end observability platform that spans from the core infrastructure to the final user experience. Collectively, XSIAM and Observability now represent over $1 billion in ARR, a remarkable achievement for data-intensive platforms that were not part of our portfolio just a few years ago.

The cornerstone of our success throughout my tenure at Palo Alto Networks has been our ability to identify premier technology and world-class talent and seamlessly integrate them into our culture. While the complexity of our integration effort naturally increased the scale of this year's acquisitions, the result has been extraordinary.

In Q4, the success was most evident in our performance with CyberArk, or now called Idira. Just two quarters after finalizing our largest acquisition to-date, we are accelerating growth while capturing synergies ahead of schedule, a rare feat that demonstrates the power of our integration engine.

These results are a testament to the deep collaboration with our new colleagues. From a go-to-market perspective, our joint efforts yielded over 400 shared leads, driving more than 200 net new logos from our installed base. We are also seeing a significant move towards larger commitments with $5 million plus TCV deals up 50% year-over-year in the fourth quarter.

Yet, the most significant challenge and opportunity remains the rise of agentic AI. By definition, an agent possesses agency, necessitating a machine identity with the precise context and permissions required to execute its workflow. As enterprises deploy thousands of these autonomous entities, many remain outside of formal governance, often lacking properly scoped permissions.

This summer served as a wake-up call as rogue agents compromised environments at several frontier AI labs. In one notable instance, an agent escaped its sandbox and exploited system vulnerabilities because its access had never been properly restricted. At its core, this represents a fundamental identity crisis for the enterprise. This is the strategic imperative behind our Idira platform.

Idira extends sophisticated identity security and privilege controls to AI agents, ensuring every machine action is authorized, scoped, and fully auditable. As we integrate these agentic controls with our AI gateway into Prisma AIRS, we are empowering organizations to enforce security policies and maintain defense in real-time.

Fiscal 2026 was a transformative year for Palo Alto Networks and the broader industry. We remain convinced that the AI tailwinds catalyzing cybersecurity demand will only intensify as we look towards the future. First, the global AI infrastructure buildout is drawing trillions in investment. We anticipate more capital expenditure in the next five years than the preceding two decades.

This massive expansion is fueled by demand that continues to outstrip supply. For AI to deliver on this promise, both traffic and data volume must scale, and as they do, every bit requires inspection and every byte requires observability. This surge in critical infrastructure is a permanent tailwind for cybersecurity, a trend already manifesting in the accelerated momentum of our Network Security and Observability businesses this year.

Second is the strategic imperative to transition towards real-time defense. With cyberattacks now operating at machine speed, fragmented legacy tools are no longer viable. There is approximately $1 trillion of global cybersecurity debt that must be modernized to defend against automated threats. Because AI operates instantaneously, this modernization must occur on unified platforms. Platformization is the only solution for real-time defense, ensuring that telemetry and policy are harmonized across every control point. We're still in the early chapters of this structural change.

Third, AI has inaugurated a fundamentally new market for cybersecurity. The rise of autonomous agents will dramatically expand the network surface area that requires fortification. Robust governance and security guardrails for AI have shifted from optional features to essential enterprise requirements. While this market is

evolving rapidly, we believe the future belongs to architectures providing end-to-end controls, a vision we're delivering through Prisma AIRS.

Lastly, I do want to mention in breaking news, we closed our acquisition of Console today. Console brings an AI-first approach to product development in the IT and security operations space. Andrei and his team are going to work as part of our Cortex effort to agentify our capabilities and drive us faster into the AI era. I want to welcome both the Embrace and Console teams, acquisitions we closed this quarter, to Palo Alto Networks.

As we move into fiscal 2027 with significant momentum, we understand that our continued leadership must be earned through disciplined execution every quarter. I want to express my gratitude to our employees for their performance during this milestone year and to our customers for their enduring partnership.

With that, let me hand it over to Dipak.

Dipak Golechha

Chief Financial Officer, Palo Alto Networks, Inc.

Thank you, Nikesh. And good afternoon, everyone. We delivered a strong close to a record year, driven by the broad-based strength across our platforms and the early success of our integration efforts. Our teams executed with discipline, and we exceeded guidance across every metric.

Before walking through the details, please note that I'll be speaking to our results both on a reported and a pro forma basis to provide a normalized growth comparison where applicable. All growth percentages will be on a year-over-year basis unless stated otherwise.

Starting with the top line, Q4 RPO exceeded $20 billion for the first time, ending the year at $21.2 billion, up 34%. Our bookings growth accelerated for the second consecutive quarter on a pro forma basis, driven by the success of our platformization strategy. Current RPO reached $9.3 billion, also up 34%, as contract durations remained steady year-over-year.

We also delivered a record result in NGS ARR, which reached $9.1 billion in Q4, up 63%. As Nikesh highlighted, most notable was the nearly $1 billion of net new NGS ARR in Q4, which almost doubled year-on-year and is a milestone that only a select category of technology companies have ever achieved.

I still recall my first quarter as CFO in Q3 of fiscal 2021 when we surpassed $970 million in total NGS ARR. We've now added approximately that amount in a single quarter. That's a testament to the multiple growth drivers in our business. Five years ago, SASE was still in its infancy and XSIAM had not yet launched. Today, those have either surpassed or are approaching a $1 billion ARR businesses.

To provide more visibility into our growth drivers, we're introducing new revenue disclosure by platform, as I previewed last quarter. Those three platforms are Network & AI Security, Cortex, and Idira. We provided historical periods as well as product composition for these platforms in the appendix of our earnings presentation published on our website.

Before diving into our revenue by platform, please note that Network & AI Security includes the Certificate LifeCycle Management business we acquired with CyberArk, which has since been rebranded to Next-Generation Trust Security, or NGTS. NGTS contributed approximately $85 million to Network & AI Security revenue in fiscal year 2026.

Additionally, the revenue by platform I will discuss excludes certain items like professional services, which are reported in the category titled Other, as shown in the earnings presentation appendix.

Let's start with Network & AI Security. Our revenue here grew 17% for the full fiscal 2026, reaching $8.35 billion in revenue. We continue to deliver above-market and double-digit growth in Network Security, which speaks to our strong competitive position and the large market opportunity still ahead of us in our largest platform. As an example, we continue to gain share in SASE, with bookings and ARR growing well ahead of the overall market.

Our Software Firewall business accelerated once again, reaching 29% ARR growth in Q4 and Prisma AIRS surpassed $100 million in ARR within its first year of general availability. Finally, we had another strong quarter in our Hardware Firewall business, driven by the adoption of our latest Gen 5 appliances.

Turning to Cortex, which includes our Security Operations and Observability platform, revenue grew 25% in fiscal year 2026 to $1.92 billion in revenue. As noted earlier, XSIAM continues to be a key driver of Cortex, with ARR growing 70% in Q4.

On the Observability side, our ARR surpassed $500 million and more than doubled since we closed the acquisition of Chronosphere in Q2. Keep in mind, and as we noted last quarter, our Q4 net new ARR includes a nine-figure benefit from a large LLM customer migrating to Chronosphere from an incumbent vendor.

Lastly, we have Idira, which consists of our Identity Security platform from the CyberArk acquisition, closed in early fiscal Q3. As noted earlier, Idira excludes revenue from the Certificate LifeCycle Management acquired from CyberArk.

On a pro forma basis, Idira revenue reached $1.26 billion in fiscal year 2026 and grew 21%. Our bookings grew faster than revenue in Q4, which is a testament to our early integration success and go-to-market collaboration. In total, our revenue grew 34% to $3.41 billion in the fourth quarter and for the full fiscal year, revenue reached $11.5 billion, up 24% year-over-year. From a geographic perspective, we delivered robust growth across all of our regions. The Americas was up 33% year-over-year, EMEA was up 39% year-over-year, and JPAC was up 34% year-over-year.

Moving down the P&L, total gross margin in Q4 was 74.8%, down 100 basis points year-over-year. For the full fiscal year, gross margin was 75.8%, down 60 basis points year-over-year. This decline reflects a mix shift towards our faster-growing SaaS offerings, which continue to scale with our platforms and have yet to reach their gross margin maturity. Looking ahead, the growing majority of revenue is cloud and SaaS, and we anticipate that mix shift will drive our cloud hosting costs faster than total revenue in fiscal year 2027.

Turning to the supply chain, we expect rising commodity costs to persist in our hardware business, particularly as it relates to memory and storage. As a reminder, while we're pleased with the strength that we're seeing in our hardware demand, revenue from hardware represents approximately 10% of the total company. We continue to manage our component cost exposure through our strategic supplier relationships and selective pricing actions across our portfolio of hardware products. Ultimately, our primary focus remains on optimizing the business for total operating income and margin, and this focus was reflected in our Q4 results and our full-year results.

Q4 non-GAAP operating margin came in at 29.6%, and for the full fiscal year, we achieved operating margin of 29.2%, an increase of 40 basis points year-over-year. This annual expansion is particularly notable as it includes a partial year of our largest acquisitions, which operated at much lower operating margins at stand-alone entities. We're making excellent progress on this front. Regarding CyberArk synergies, our integration synergy targets

remain three to six months ahead of plan. Looking ahead to fiscal year 2027, we anticipate higher cost of goods sold will be more than offset by continued operating leverage as we scale efficiently and deliver on M&A synergies.

Our focus on operating leverage drove Q4 non-GAAP EPS of $1.02, exceeding the high end of our guided range by $0.04. Adjusted free cash flow for the fourth quarter reached $1.29 billion, growing 35% year-over-year. For the full fiscal year 2026, adjusted free cash flow was $4.41 billion, delivering a margin of 38.4%, an increase of 40 basis points year-over-year. As a result of our strong free cash flow generation, we ended fiscal 2026 with a robust balance sheet, including $7.9 billion in cash, cash equivalents, and short-term investments (sic) [short-term investments and long-term investments].

Stepping back, over the past three years, we've proven our ability to deliver durable and profitable growth. Our execution has driven over 500 basis points of operating margin expansion. We've achieved this whilst capturing market share across new categories, driven by our industry-leading R&D investment. Our operating leverage has also translated directly to cash flow. Adjusted free cash flow margin has been 38% or better in each of the last four years, and we sustained the strong cash flow generation even while absorbing the impacts of large M&A and as our customers moved increasingly from multiyear to annual billing.

This track record of scaling profitably is the bedrock of our financial model. It provides us with the ability to neutralize potential cost headwinds while simultaneously fueling our innovation engine, our ultimate competitive advantage and the catalyst for our customers' platformization journeys. Looking ahead, we continue to have increasing visibility into our free cash flow. This is being driven by a combination of steady operating margin expansion as well as a smooth transition to deferred or annual billing in our core business.

To provide some context, annual billings increased significantly from 6% of bookings in fiscal 2020 to 27% in fiscal 2025. Now we're seeing a steadier rise with the percentage of annual billings having increased by low-single digits year-over-year in fiscal 2026 to about 30% of total bookings. With this structural transition now largely stabilized, we have highly predictable compounding cash engine going forward. This cash flow visibility, paired with our continued focus on margin expansion and durable double-digit bookings growth, reinforces our confidence in achieving our 40% free cash flow margin target in fiscal 2028.

Before we turn to guidance, I also want to step back and frame the growth opportunity ahead. As I mentioned earlier, our industry-leading R&D investment over the years has fueled our innovation engine and expanded our market opportunity into new categories. That ongoing commitment has earned us leadership recognition in nearly every major category that we operate in. What began predominantly as a stand-alone firewall business is now a platform with multiple billion-dollar ARR businesses and several more approaching that milestone.

We continue to remain under-penetrated against our total addressable market of $340 billion by 2030. We believe that AI will only expand our opportunity whilst reinforcing the need for platformization and real-time cyber defense. This puts us on track to achieve our target of $20 billion in NGS ARR by fiscal year 2030.

With that long-term framework in mind, let's turn to our Q1 and our fiscal year 2027 guidance. Note that our recently closed acquisitions of Console and Embrace are immaterial to our fiscal year 2027 guidance. For the fiscal first quarter 2027, we expect – for Q1, we expect NGS ARR of $9.54 billion to $9.56 billion, or 63% growth. We expect RPO of $20.8 billion to $20.9 billion, or 34% to 35% growth. And we expect revenue of $3.3 billion to $3.31 billion, or 33% to 34% growth. Fully diluted share count of 837 million to 844 million shares and diluted non-GAAP EPS to be in the range of $0.96 to $0.98 per share.

For the fiscal year 2027, we expect NGS ARR of $11.075 billion to $11.175 billion or 22% to 23% growth, we expect RPO of $25.2 billion to $25.4 billion or 19% to 20% growth, and we expect revenue of $14.1 billion to $14.2 billion or 23% to 24% growth. We're guiding operating margin of 29.5% and diluted non-GAAP EPS to be in the range of $4.16 to $4.19 per share. Fully diluted share count of 844 million to 847 million shares and adjusted free cash flow margin of 38%.

We've included our typical modeling points in the appendix of our presentation for your review, but I would like to point out a few things. First, as previously mentioned, our fiscal year 2026 net new NGS ARR included a nine-figure benefit from a large LLM customer migrating to Chronosphere from an incumbent provider. Our outlook assumes the tail end of this migration will last through Q1 of fiscal 2027 and that the net new ARR contribution from this migration will be less than what was added in Q4. This will impact the seasonality of the net new NGS ARR in fiscal 2027, making Q1 larger than normal. We expect 60% to 61% of the net new NGS ARR to fall in the second half of fiscal year 2027.

Second, while we do not intend to give revenue guidance by platform, we are providing initial modeling points to help you establish the revenue growth trajectory for each of the platforms within the context of our total company guidance. For fiscal year 2027, we expect Network & AI Security revenue growth of low-double-digits year-over-year, we expect Cortex revenue up approximately 30% year-over-year; and we expect Idira revenue of approximately $1.5 billion, representing pro forma growth of high-teens to 20% year-over-year.

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

QUESTION AND ANSWER SECTION

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

A

Okay. Thank you, Dipak. Please ensure that only one question is asked by each analyst. First question will be Rob Owens from Piper Sandler, followed by Brian Essex from JPMorgan.

Rob D. Owens
Analyst, Piper Sandler & Co.

Q

Great. Thank you, Hamza. And thank you, guys, for taking my question. Nikesh, your prepared remarks spoke to a lot of the tailwinds that you guys are seeing across cyber right now. And I think that was evidenced in your bookings strength and you mentioned the second straight quarter of acceleration, but this has been uneven throughout the environment, and obviously scaled players and players with breadth of coverage really has mattered here.

So, to that end, as you look at the new fiscal year, how are you thinking about M&A, how are you thinking about something else that could be transformational to Palo Alto, just given that the market is shifting so quickly? And while you have had an ability to take advantage of it, given what you've done in the past, what are you contemplating moving forward? Thanks.

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

A

Rob, thank you for your question. I'll just send you the names of the company so it makes it easier. I don't have to answer them in such detail.

Rob D. Owens
Analyst, Piper Sandler & Co.

Q

We'd appreciate...

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

A

You'd appreciate that, right?

Rob D. Owens
Analyst, Piper Sandler & Co.

Q

We'd appreciate that.

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

A

Now, look, as I've always maintained that M&A is not a strategy, M&A is a consequence of stuff that we do from a product development perspective. To give you a sense, if you – I talked about the three major pivots we've seen in AI already in the last seven months. You've seen people go from LLMs to agents to now open-weight models. And every one of these technological shifts on the customer side obviously requires a slightly different security architecture. How do you protect these agents, how do you ensure that open-weight models are protected, agents don't go rogue? And obviously, we have a point of view internally, and we're building towards that from a product development perspective, but sometimes you can get caught flat-footed because you're going down one path and suddenly the market shifts elsewhere.

This is where I – we have the privilege of looking at the entire cybersecurity landscape and seeing 40 or 50 companies that have been funded in this category, and then you suddenly realize that some other company had the strategy right, and that's when you step in and make an acquisition. So, the acquisition happens because they've got a technology trend right and we'd rather embrace it quickly and get on that, so our customers can have that capability much faster. Because, honestly, as you can see, after Mythos, what has happened is customers are willing to experiment with a lot of AI implementations, but before they deploy, they want to ensure a robust security harness around it. The most sort of common questions we get are, what do I do about the vulnerabilities that Mythos is going to find in my environment, how do I solve it today and how do I follow it for the long term, or what happens if we deploy agents and our agents go rogue, how do we make sure our agent doesn't go running to Hugging Face?

Rob D. Owens
Analyst, Piper Sandler & Co.

Q

Excellent. Thank you.

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

A

Thanks, Rob.

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

A

All right. Thank you, Rob.

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

I'll keep your request, and I'll send you the company's name as soon as I buy it.

Rob D. Owens
Analyst, Piper Sandler & Co.
Q

I appreciate that.

Hamza Fodderwala
Senior Vice President-Investor Relations & Strategic Finance, Palo Alto Networks, Inc.
A

All right. Thanks for the question, Rob. Next, we have, Brian Essex from JPMorgan, followed by Saket Kalia from Barclays.

Brian Essex
Analyst, JPMorgan Securities LLC
Q

Hey. Great. Thanks for taking the question. Nikesh, it's great to see the acceleration in CyberArk performance, in only 200 net new logos from the Palo Alto installed base. Would love to get a sense of what those conversations are like, how big are those deals relative to the rest of the CyberArk platform? And you still have a substantial amount of your installed base. I think a lot of people focus on the cost synergies. They forget about the revenue synergies. How much penetration do you think you can get into your installed base with the CyberArk platform? Thank you.

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

Look, I'm really excited about CyberArk. I think, if you look at both ends, and you rightfully articulated, we have been able to really hit the ground running on the cost synergies side. You've seen that our margin is reverting back to what our stand-alone margin was in just about two quarters. And we think we'll be at a stable point coming into the next quarter. So, to be able to transform a large company like CyberArk in nine months and get their margins up by 1,000 basis points or more is already good work on the cost side. But like you said, we didn't buy it because we had cost synergy. We bought it because we felt there's a need in the market for identity security, and this was an inflection point.

I think the phase one, from our perspective, was don't break it, accelerate their momentum. And you've seen we've been able to do that. We've just hired a new leader last quarter, Sonny Singh. He's right now at our sales conference in Asia, rallying the troops in CyberArk. The team has taken really well to joining Palo Alto. I think there's been phenomenal collaboration between the two teams. I'm excited. We just launched a new product called Modern PAM. So, CyberArk was in traditional PAM. Modern PAM is a expansion category for PAM, something they hadn't spent a lot of time on before. The product team at CyberArk has been – or, Idira now, I should say, has been amazing at being able to embrace it. That product is generally available now. We expect to try and upgrade all of the existing traditional PAM customers to that.

So, there's a lot of activities we have going on in both on the upsell and expansion side as well as the net new sales side. So, as long as we can run at a faster growth rate than CyberArk ran individually – independently and expand the margin by 1,100 basis points, I think that's a phenomenal acquisition for us, not to mention that they have a pole position in being able to help with nonhuman identities and agents going forward because that is a whole new field where there is no established leader.

Brian Essex
Analyst, JPMorgan Securities LLC

Q

Got it. Very helpful. Thank you.

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

A

All right. Thank you, Brian. Next, we have Saket Kalia from Barclays, followed by Fatima Boolani from Citi.

Saket Kalia
Analyst, Barclays Capital, Inc.

Q

Okay. Great. Hey, guys. Thanks for taking my questions. Great finish to the year. Nikesh, maybe for you. You've said that Mythos isn't a moment, but it's rather at the beginning. And so, maybe the question here is, how are you seeing buying behavior change as the AI threat becomes the new normal? And what I mean by that is, do you see more of a willingness to platformize? Do you see more pipeline growth than you would expect? Do you see more appreciation for value, less sensitivity in pricing? I guess I'm just curious if you can translate this new beginning to some of the deal dynamics that you saw in the quarter or the last couple of quarters?

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

A

Please make sure the sweets show up at Hamza's house a week before. Otherwise, you won't get your first spot to ask questions in the future. So, in terms of the momentum, look, I did say Mythos is the beginning because what is happening is I've been strived for eight years to try and get CEOs interested in cybersecurity, I couldn't, but Dario did a phenomenal job by having Mythos because every CEO now wants to talk about what does this mean to us, how do we get access to it, how do we test ourselves from a vulnerability perspective?

But they're wise. They sit there and say, listen, I get it that this is the new normal. People will be able to find vulnerabilities much faster. How do I solve this problem in the long term? That's really where the conversation starts about the only way to solve this problem in the long term is if something escapes past your perimeter, you've got to find it quickly and shut it down. That talks about modernizing their cyber estate, that talks about platformization, that talks about having an AI-driven SOC. So, that's why we've been able to have so many conversations around the modernization of infrastructure.

And every conversation is not about fragmenting their estate and buying yet more smaller vendors. It's more about finding a consolidated way of sort of standardizing on a platform, evaluating a platform. I think this is a big tailwind for the larger players in the sector. I don't think this is a moment where – you will see obviously start-ups with some unique products, niche products, which they are able to bring to market faster, which customers will use in the interim, but I think this is definitely a long-term, I'd say, duration-changing trajectory change to our growth rate.

Because, you think about it, open-source models are now already able to compete with the capabilities of Mythos, and this thing is going to get better, not worse. If that happens and this capability becomes commonplace, we have a short window by when to get all the cybersecurity technical debt, which hasn't been paid over the many years, back up to the mark. I suspect there will be some major breaches over the coming years because customers have not been able to get their transformation act in place. And that's generally going to be a tailwind for all of us in this space.

Saket Kalia
Analyst, Barclays Capital, Inc.

Q

Very helpful. Thank you.

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

A

Okay. Thank you, Saket. Next, we have Fatima Boolani from Citi, followed by Matt Hedberg from RBC.

Fatima Boolani
Analyst, Citigroup Global Markets, Inc.

Q

Thank you for taking my question. Nikesh, you brought up this concept of technical debt. So, I wanted to zoom out and ask you a question in the context of something you announced earlier this week or a couple of weeks ago, Frontier AI Critical Defense. So, one thing we haven't necessarily heard you talk about is this notion of operational technology and the use case here potentially gaining critical mass, and especially in the context of your own platformization strategy.

So, now that we know what the models are capable of in terms of insane vulnerability chaining against a part of your technical environment that has historically been underinvested in, again, with a lot of technical debt, what are some of the gating factors here still for you to be able to accelerate wallet capture? And then, relatedly, how does that cooperation versus competition continuum with some of the frontier lab partners that you have get expressed in this market opportunity with OT that seems like it would be ripe for more capture?

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

A

Fatima, a lot of questions in there. Look, first and foremost, I think nine months ago, we were all guilty and convicted of near death as cybersecurity and software because frontier AI was going to eat all of our lunch and breakfast and dinner. Clearly, in the last six to nine months, it's become apparent that that's not happening. We're all going to be enjoying this feast together. And we've seen both OpenAI and Anthropic and even Google come to the table in terms of partnerships. We have early access to these models. We're able to test them. We're able to test their cybersecurity capabilities.

As I said in my prepared remarks, we were the first or are the first commercial partner allowed to use Mythos as part of our testing harness. We already use OpenAI 5.6 as part of our testing harness. We are able to bring multiple models to customers. Because the customers are quickly disenchanted from this notion of finding more vulnerabilities. They want to know, what do I do about them? The last thing they want is more security problems. They have enough already. So, the conversation is quickly shifting from, what do I do about this? And in that conversation is where the need for platforms, as I was mentioning earlier, comes up.

In terms of OT specifically, I think the challenge is even more pronounced because OT is hard to patch. Even if you found a vulnerability in an OT instance or deployment, imagine patching an oil rig out in the ocean or imagine patching a bunch of technology which does not have remote access, cannot be remotely patched, you'd have to go there and fix it. The good news is – Lee is not here this week, so I'm going to do Lee right now. So, we have actually built a capability where we can build signatures for OT vulnerabilities and open-source vulnerabilities and deploy them in under four hours.

So, we can find an open-source vulnerability and OT vulnerability, deploy the fix in four hours, and propagate that to our software and hardware firewalls so that it'll stop the bad actors in their tracks, which is a far cry from the

current standard of 55 days. It takes 55 days to patch open-source vulnerabilities or OT vulnerabilities in the wild. This will allow our customers to have the ability to block the bad actors for any network-related OT or open-source vulnerability in under four hours.

So, it's a good thing you asked me what the gating factor was. The gating factor is really the customers taking the time to understand what major changes do they need to make, doing POCs, assessing what the environment looks like, thinking about who they want to deploy, then eventually getting down to deployment. This is not something customers are – they take their time to go do the deployment.

And that's why I think it's a long-term tailwind, and you will start seeing that in constant sort of over-performance in the industry on a quarterly basis. But it's not going to be coding agent-style ARRs that we're seeing in the AI space, which I'm envious of, but yeah...

Fatima Boolani Q
Analyst, Citigroup Global Markets, Inc.

Good Lee answer, but not good sideburns, Lee sideburns.

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

Well, that's easy to fix.

Hamza Fodderwala A
Senior Vice President-Investor Relations & Strategic Finance, Palo Alto Networks, Inc.

Okay. Thank you...

Fatima Boolani Q
Analyst, Citigroup Global Markets, Inc.

Thanks, Nikesh.

Hamza Fodderwala A
Senior Vice President-Investor Relations & Strategic Finance, Palo Alto Networks, Inc.

...Fatima, for the questions. Next, we have Matt Hedberg from RBC, followed by Michael Turrin from Wells Fargo.

Matthew Hedberg Q
Analyst, RBC Capital Markets LLC

Thanks, Hamza. Nikesh, you guys have a long-standing vision of being the number one vendor in a category. I mean, you don't enter a market unless you think you can be the share leader. And so, I guess, putting Lee's hat on again, you've had a lot of success, obviously, in observability, with stand-alone Chronosphere, you added Embrace, Synthetics – or you developed Synthetics. Where are you from a functionality perspective now versus some of the sort of the historic market leaders there? And how much of this is share shift versus just like this market's just getting bigger with AI and we think we can take a lion's share of it?

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

Well, look, the premise of Chronosphere has been that it was designed for the AI era. It is a net new technology. The premise of Chronosphere is that because of the large volumes of data that are being sort of spit out in the

observability space, it is designed as an architecture that allowed you to have a lower total cost of ownership. So, Chronosphere is on average 30% or 40% cheaper than any of the leading incumbent observability solutions out there.

From a parity of capability perspective, we started off being very good from an AI-native perspective, from traces, logs, and metrics. So, a majority of Chronosphere's customers are AI-native customers, including a very large frontier AI lab. With the absorption of Embrace and the development of Synthetics, that will put us at par with some of the leading players on a cross sort of capability perspective, which allows us to go after the enterprise space.

So, that'll allow all the Palo Alto sellers to start selling. For now, we're restricting Chronosphere just to AI-native sales because it's where it's most suited. But I expect in the next six months, we'll get to a point where Chronosphere will be a competitive product in its category vis-à-vis other enterprise players, and then we have both an AI-first capability as well as a cost advantage.

So, that should allow us over time, as the space normalizes, to have a multibillion-dollar ARR business. Very excited. We bought it when it was $85 million (sic) [$200 million] ARR. It's already crossed $0.5 billion in ARR. We can clearly see a line of sight for that to keep getting bigger over the next few quarters, and then hopefully address the enterprise market with it as well.

Because, remember, for us to reach our aspirations of a bigger business, we need to have multiple multibillion-dollar ARR businesses. Observability is such a TAM, SIEM is such a TAM, and obviously, our Network Security business and Identity business are similar TAMs.

Matthew Hedberg
Analyst, RBC Capital Markets LLC

Q

Thank you.

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

A

Thank you, Matt. Next, we have Michael Turrin from Wells Fargo, followed by Gray Powell from BTIG.

Michael Turrin
Analyst, Wells Fargo Securities LLC

Q

Thanks very much for taking the question. Great close to the year. Maybe just on the initial fiscal 2027 guide, I'm curious how you approach that exercise given the inflection points taking shape across cyber. You mentioned three major AI inflections you've seen form. We're still early in the overall 2027 cybersecurity budget discussion. So, maybe just walk us through what you're assuming as a baseline and any key drivers of upside you see on the horizon we should focus in on as well? Thank you.

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

A

Michael, if we take the guidance very thoughtfully and we look at where you are from a consensus perspective, we make sure we look at the underlying business plans of our businesses, evaluate if we are able – going to be able to meet, beat, or exceed your consensus. We're delighted to see that we expect with our execution and the tailwinds, we are going to be able to exceed your consensus. And that's how we guide.

Michael Turrin
Analyst, Wells Fargo Securities LLC

Q

It's very clear. We look forward to it.

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

A

Go ahead, Dipak.

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

A

Yeah. No, I think, Michael, look, we do look at a lot of different inputs. If I just look at a number of the different trends, we will look at what's happening to pipeline, are we seeing traction, do we see a trend in terms of what's going on with some of the new areas that we have? We take all of that, ingest it all, look at the resource requirement required for territory planning, et cetera, et cetera. And that's effectively how we do it. It's a pretty well-established, world-class process. I wouldn't say much has changed from a process point of view in the last five, six years that I've been here as the CFO. And I think we've been pretty transparent. And there have been a number of inflection points that we've been able to kind of like capture within our forecast criteria.

Michael Turrin
Analyst, Wells Fargo Securities LLC

Q

Thank you.

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

A

All right. Thank you, Michael. Next, we have Gray Powell from BTIG, followed by Meta Marshall from Morgan Stanley.

Gray Powell
Analyst, BTIG LLC

Q

Great. Thanks for taking the questions and congratulations on the really strong results. So, I just want to make sure that I was looking at something correctly. I think last quarter, you called out $200 million in competitive SASE displacements for the last nine months. This quarter, that number jumped to $450 million. So, I just want to make sure that those are comparable statistics, because if so, well, you had a really big Q4. Either way, the numbers are impressive. What's driving the improved pace of displacements and just overall strength in SASE relative to peers?

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

A

Gray, I think the number is $400 million, if I remember correctly. $450 million? Okay, it's $450 million. Good. Well, clearly, we had a good Q4. That's evident in our numbers. So, yes, we did have a good Q4. Look, the displacement is a consequence of two events. One, when SASE as a category came about early, it was a very Internet-driven phenomenon. It was Internet access-driven. But COVID changed all of that. When we hit the COVID mark, people wanted sort of access consistently both to the private access as well as Internet access, which is where we come from. We come from the private access space. And obviously, our product on the Internet access space is now at par or far exceeds the competitive landscape we have in front of us.

It's really the sort of integration of SASE with SD-WAN, which we were early in. We were the first player to go acquire CloudGenix, integrate that to a SASE fabric. Having our SASE fabric be consistent with our hardware and software firewall fabric allows our customers who use Palo Alto firewalls to actually gravitate towards our SASE solution as opposed to elsewhere. And not just that, it also makes it an easier choice if they're looking to consolidate and have one platform because they already are using our consoles, our Strata Cloud Manager, our services for the hardware and software firewall use case, then it doesn't feel like a big sort of change or to go like adopt us on the SASE front as well because they already also have our agents in many cases, which do the VPN product, is now a consistent agent of SASE.

So, we've surrounded the SASE sort of incumbents with effectively a complete platform where the choice of standardization on our platform is a simpler choice for them if they choose to just replace the SASE piece, because they already have the other elements from us. So, sometimes it's that, sometimes it's just perhaps customers who want to modernize their SASE infrastructure.

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

And just to clarify, Gray, it was $200 million year-to-date at Q3, and it's $450 million for the full year.

Gray Powell Q
Analyst, BTIG LLC

All right. So, that's a pretty big number for Q4. Thank you. That all makes a lot of sense.

Hamza Fodderwala A
Senior Vice President-Investor Relations & Strategic Finance, Palo Alto Networks, Inc.

Okay. Next, we have Meta Marshall from Morgan Stanley, and our last question will be Brad Zelnick from Deutsche Bank.

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

Great. Thanks. Nikesh, you were mentioning kind of this addressing of the $1 trillion of technical debt. Platforms can help enterprises pay for that in some ways. But just how do you think either about ways that you can help them in terms of professional services, investment, or other things that can help from just speeding up the amount of technical debt they can address in a compressed period of time? Thanks.

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

So, as you know, Meta, a few years ago, when we launched the platformization strategy, we've had very clear models in the market where we were willing to take staggered payment or align their contracts or deploy before their existing vendor has to be replaced to drive faster platformization. So, we make all that available.

Honestly, the constraint that you always run into is the customers always have a full deck. They're already working on a series of things that they would like to get done in their enterprise. And today, with AI, there's a very large contingent of AI transformation that's out there. People want to transform customer support. They want to go do coding on an aggressive basis. They want to deploy LLMs. So, this is yet another priority that must be managed in the context of that overall priority.

So, it's just a balance the customers strike. That's why they don't go whole hog and say, let's go replace everything tomorrow. They do sit down and say, let's have a more cohesive and intelligent transformation plan. As a transformation plan, it's going to take five years, which is too long, we got to get it done sooner.

So, you typically end up in the one, two, three range, but it's not something that gets done in one quarter. And they want to sort of crawl, walk, run. They want to get some stuff done as other vendors sort of fall off their sort of end-of-life periods or their contracts are up for renewal.

So, all I can say is the desire to standardize or platformize on larger vendors where products are at par or better than the state-of-the-art of the market is becoming more and more of a trend, and that's generally in our favor.

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

Q

Great. Thanks.

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

A

Right. Thank you, Meta. And last, but certainly not least, we have Brad Zelnick from Deutsche Bank.

Brad Zelnick
Analyst, Deutsche Bank Securities, Inc.

Q

Wonderful. Thanks very much, Hamza. Nice to see everybody. Nikesh, you have strong credibility doing M&A at this point, and today's Console acquisition seems directionally consistent with moving closer to autonomous security operations. And I can ask the simple why Console, but if you fast-forward five years and Palo Alto has succeeded beyond your wildest expectations, what's the most valuable activity that customers have completely stopped doing themselves because Palo Alto Networks is doing it for them?

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

A

It's a great question, Brad. I think that's why I know why Hamza saves you for last. So, look, if you believe that we're going to spend $5 trillion of CapEx in the next five years building data centers and AI capability, I have to believe that AI is going to be adding tremendous value to our lives in the enterprise space. Otherwise, it makes no sense to deploy $5 trillion in the ground.

So, I'm an optimist and believe that we will be using a lot of AI to do a lot of agentic tasks. And if that's true, cybersecurity has to become less manual and more agentic and more done by us than the customers themselves, because the bad actors will be using AI from their angle, which means we have to make sure our customers are as agentified or AI-fied as the bad actors are.

Now, that is not possible as you're discovering in every industry category. You cannot deploy AI effectively until you have the right data in place, the right training data, the right data. You have to break the silos and have things, talk to each other. Now, that leads itself towards a cohesive, unified data lake of some sort, whether it is an enterprise IT data lake, observability data lake, a security data lake.

If you see strategically where we have been pivoting the business over the last two or three years is we're a very data-first company. We ingest a lot of data in XDR. We ingest 19 petabytes a day in the SIEM product already,

and we have just barely north of 1,000 customers. We have observability data, which is now the data of an entire frontier LLM that is being ingested to provide them observability.

So, we are becoming a data-oriented, AI-first cybersecurity company. Our aspiration is to reduce the amount of human intervention in the act of detection, prevention, and remediation in the cyberspace. So, if you were to ask me what's that North Star, that's our North Star. The question is, how do we get there? And that's where the whole company is focused on trying to get there.

So, in five years from now, if we were far exceeding our expectations of ourselves, I would be able to walk in to a company and say, you want to replace X? Guess what, I have agents that can understand your deployment. My agents will replace that product. And I can do that in under a week.

And when I deploy my product, you'll need a lot less people, and our products will actually just look for validation from you and get the task done, without having you to get into the nits and grits of how to configure things and what policies to write, because we've seen that across thousands of instances and we can bring that intellectual knowledge to bear.

Today, if we look at enterprise products, every enterprise product starts dumb for the next customer despite being deployed for 100,000 customers. I think AI gives us the opportunity of learning from the multiple deployments we do and the multiple customers we have and show up more intelligent for the next customer every time. And that's the aspiration we have.

Brad Zelnick Q
Analyst, Deutsche Bank Securities, Inc.

Makes sense. Thank you.

Hamza Fodderwala A
Senior Vice President-Investor Relations & Strategic Finance, Palo Alto Networks, Inc.

Thank you...

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

Thanks, Brad.

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

All right. That concludes the Q&A portion of the call. I'll hand it back to Nikesh for any closing remarks.

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

I just want to take the opportunity to once again thank all of you guys for being here. Thank our customers, our shareholders, and all of our employees for what was a spectacular FY 2026 for all of us at Palo Alto Networks.