Omdia
by informa techtarget

The State of Network Security in the AI Era

How Palo Alto Networks Software Firewalls Secure Multi-Cloud and AI Workloads

Author: John Grady, Chief Analyst

October 2026

In partnership with:

paloalto NETWORKS

© 2026 TechTarget, Inc. All rights reserved. Unauthorized reproduction prohibited.

This Omdia White Paper was commissioned by Palo Alto Networks and is distributed under license from TechTarget, Inc.

Contents

Executive summary 3

Foundational distributed infrastructure challenges can be overlooked..... 4

Why network security needs to evolve to address the modern multicloud, AI-centric world 5

Key requirements for protecting modern distributed infrastructure and why traditional solutions fall short 7

Modernizing firewall architectures with Palo Alto Networks Software Firewalls..... 9

Conclusions 13

© 2026 TechTarget, Inc. All rights reserved. Unauthorized reproduction prohibited.

This Omdia White Paper was commissioned by Palo Alto Networks and is distributed under license from TechTarget, Inc.

Executive summary

Security teams have an increasingly difficult charter: protect highly distributed and diverse environments without stifling innovation. The rapid onset of AI only complicates what was already an onerous mandate. Network security tools have historically been siloed across different locations and workloads, which no longer scales. Organizations need to rethink their network security approach and look to unify controls to turn the network into a single control point that can offer consistent protections across locations, architectures, and use cases. Palo Alto Networks Software Firewall platform helps security teams accomplish this by offering unified operations, a highly scalable fabric to address changing traffic requirements due to AI, and granular Layer 7 microsegmentation to prevent lateral movement. The incorporation of AI runtime security helps security teams comprehensively protect their modern AI-centric, multicloud environments.

Organizations need to rethink their network security approach and look to unify controls to turn the network into a single control point that can offer consistent protections across locations, architectures, and use cases.

© 2026 TechTarget, Inc. All rights reserved. Unauthorized reproduction prohibited.

This Omdia White Paper was commissioned by Palo Alto Networks and is distributed under license from TechTarget, Inc.

Foundational distributed infrastructure challenges can be overlooked

IT teams have been moving workloads to the cloud for well over a decade, and "born-in-the-cloud" organizations are commonplace. Yet network security teams continue to struggle with the security implications of this reality. At a high level, cloud can connote simplicity, but the truth is much more complex. In fact, nearly two-thirds (65%) of organizations said securing public cloud infrastructure is harder than securing on-premises infrastructure.[1]

Part of this is due to the fact that "cloud" is not uniform. Nearly all organizations (95%) are now multicloud, with over half (56%) supporting workloads across at least three cloud providers. In addition to location, architectures come into play. Container usage is now pervasive, with 89% of organizations deploying them in production for at least moderate use. Despite the prevalence of containers, nearly half of organizations (46%) said they cannot extend traditional network security tools into containerized environments, creating complexity as additional tools are deployed. Further, 46% also report limited visibility into east-west container traffic, resulting in blind spots attackers can exploit.

Figure 1: Foundational distributed infrastructure challenges

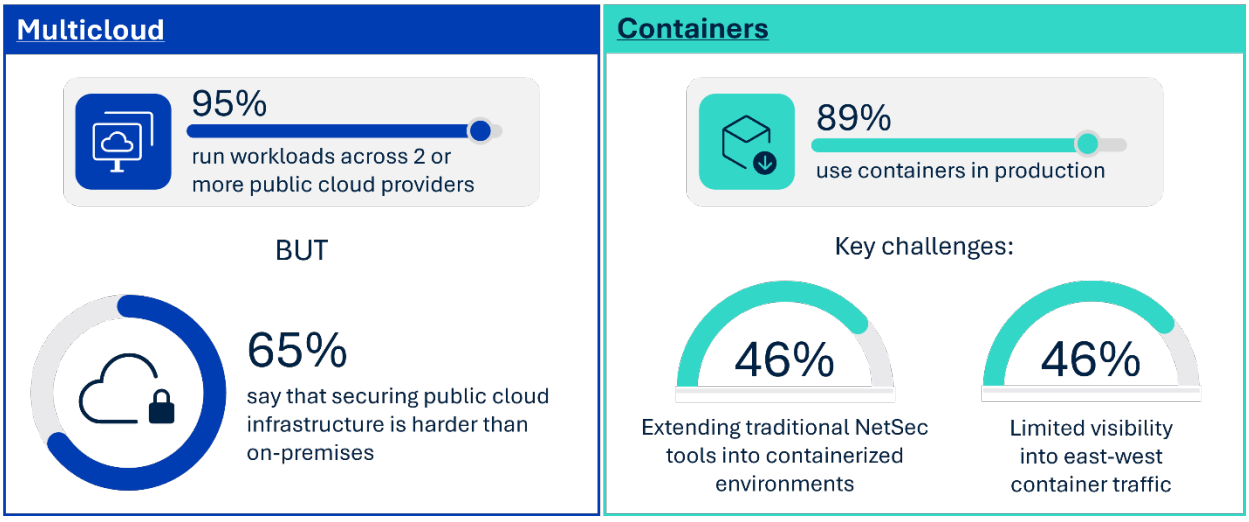

Source: Omdia

[1] Source: Omdia Research Report, *Network Security at a Crossroads: Reclaiming Relevance in Hybrid, Cloud-native Environments*, August 2026. All Omdia research references in this White Paper are from this report.

© 2026 TechTarget, Inc. All rights reserved. Unauthorized reproduction prohibited.

This Omdia White Paper was commissioned by Palo Alto Networks and is distributed under license from TechTarget, Inc.

Additionally, many enterprise organizations continue to support on-premises infrastructure, even while strategically prioritizing cloud migration. This may not be an issue on the surface, but many maintain separate network security teams, tools, and strategies between cloud and on-premises. Specifically, 26% said the practices are completely separate, while 38% said they view network security for cloud and on-premises as different but related and that there is some overlap. Only 37% cited a good deal of overlap or full unification between cloud and on-premises.

The connection paths between these different locations, clouds, and architectures often lack proper inspection when security teams and tools are focused on specific parts of the environment.

Unfortunately, attackers do not think in silos and specifically target the gaps resulting from fragmented network security practices across these different locations to gain a foothold and move laterally across the network. The connection paths between these different locations, clouds, and architectures often lack proper inspection when security teams and tools are focused on specific parts of the environment.

These issues are not hypothetical but are actively leading to successful attacks. In fact, 63% of organizations said they have experienced, or think they have experienced, an attack that exploited network security controls or configurations over the last 12 months. More than half (57%) cited attacks targeting cloud infrastructure, while 52% indicated an on-premises data center or private cloud was the target. The most common types of attacks were malware moving laterally from other parts of the environment (39%) and exploits taking advantage of known vulnerabilities (38%), highlighting the fact that attackers often take the path of least resistance and exploit holes in network security defenses.

Why network security needs to evolve to address the modern multicloud, AI-centric world

AI has only exacerbated these issues, as well as created new challenges for network security teams to address. The core issue that should be driving network security teams to reconsider their approach is that legacy firewall architectures were built for predictable, human-generated web traffic. The machine speed of AI-driven infrastructure fundamentally changes network security requirements, resulting in three primary friction points.

- **Scale**. Static hardware appliances typically come with rigid capacity limits that can throttle dynamic multicloud workloads when reached. AI workloads, while prone to

© 2026 TechTarget, Inc. All rights reserved. Unauthorized reproduction prohibited.

This Omdia White Paper was commissioned by Palo Alto Networks and is distributed under license from TechTarget, Inc.

generating traffic bursts during training and inference operations, also permanently increase baseline network capacity demands, exceeding even the 100Gbps threshold.

- **Visibility**. Traditional networks were designed primarily for north-south traffic. AI introduces significantly more east-west volume in the environment. Existing network firewalls typically stop at the cluster or VPC boundary, which creates blind spots across containerized workloads, applications, and multicloud data paths. Further, each additional AI workload introduces new data flows requiring security inspection.
- **Operational complexity**. There are existing network security tools that can address many of these issues individually. Unfortunately, this has led to firewall tool sprawl, with 91% of organizations using at least 2 providers, and 47% using three or more. Deploying, configuring, and managing these tools separately, at scale and effectively, is an impossibility for most organizations. Maintaining disparate third-party and cloud-native firewalls, along with separate on-premises and cloud network security teams, can easily lead to policy drift and misconfigurations, unnecessary administrative overhead, and, ultimately, unsustainable firewall sprawl across hybrid footprints.

With AI workloads often running inside container pods and Kubernetes clusters, maintaining Layer 7 visibility and controlling east-west pod-to-pod traffic is critical, yet lateral threat movement inside container fabrics remains invisible to traditional perimeter firewalls.

With AI workloads often running inside container pods and Kubernetes clusters, maintaining Layer 7 visibility and controlling east-west pod-to-pod traffic is critical, yet lateral threat movement inside container fabrics remains invisible to traditional perimeter firewalls. This is evidenced by the fact that attackers are already exploiting these issues (see **Figure 2**). Omdia research shows that 50% of organizations surveyed have experienced, or believe they have experienced, an attack on their generative or agentic AI applications. Nearly one-quarter (23%) indicated they have been attacked multiple times. On the one hand, it is true that attackers are typically quick to take advantage of technology shifts, so it's not a complete surprise that bad actors are already targeting AI systems. Yet the fact they are seeing success at such a high rate shows how critical it is for network security teams to catch up.

"50% of organizations have experienced, or believe they have experienced, an attack on their generative or agentic AI applications."

© 2026 TechTarget, Inc. All rights reserved. Unauthorized reproduction prohibited.

This Omdia White Paper was commissioned by Palo Alto Networks and is distributed under license from TechTarget, Inc.

Figure 2: Successful attacks on AI applications

Has your organization experienced a successful attack on any its internally built applications that use generative or agentic AI? (Percent of respondents, n=398)

Source: Omdia

Key requirements for protecting modern distributed infrastructure and why traditional solutions fall short

Despite these challenges, security professionals understand the criticality of network security for multicloud cloud and AI environments and are prioritizing consolidation around firewall platforms. In fact, 97% label firewall consolidation a critical or important priority. Yet, interestingly, these initiatives have shifted from a pure cost-saving operation to a security enhancement project, with more than half (56%) saying improved security effectiveness is a driver of their interest in firewall vendor consolidation.

"97% label firewall consolidation a critical or important priority."

More importantly, organizations are overwhelmingly bullish on using network security platforms to secure AI. Overall, 94% of security decision makers agreed that network

This Omdia White Paper was commissioned by Palo Alto Networks and is distributed under license from TechTarget, Inc.

security vendors are well positioned to protect AI applications and agentic systems, with 38% saying network security vendors are best positioned.

However, this doesn't mean that legacy firewall tools are capable of meeting the requirements needed to protect AI infrastructure at scale. To protect distributed modern infrastructure, a modern network security architecture must deliver four foundational pillars:

- **Autoscaling fabrics.** Similar to securing cloud-native applications, deployment must be automated to scale on demand. More than one-third of organizations (35%) noted that integrated load balancing and auto-scaling was an important capability in a firewall platform. Manually adding protection as new AI workloads are spun up slows the process down and creates bottlenecks. Traditional firewall architectures require intensive manual configuration, load balancing, and failover, which make addressing sudden changes in traffic difficult, if not impossible. Solutions that natively and intelligently expand as GPU demands increase help businesses take full economic advantage of their infrastructure without having to statically overprovision.

Firewall consolidation initiatives have shifted from a pure cost-saving operation to a security enhancement project.

- **Centralized multicloud governance.** Most enterprises will follow a similar model with AI that they have in the cloud: AI data, workloads, and agents deployed across multiple cloud providers, on-premises data centers and edge locations, VM- and container-based architectures, with traffic flowing both north-south and east-west. Coverage across all these elements, supported by centralized management to apply uniform policy and ensure consistent security posture across the environment, becomes critical in this situation and was noted by 39% of organizations as an important firewall platform attribute.

Firewall platforms supporting managed firewall services can remove the operational burden of infrastructure management from security teams and enable them to focus on more strategic activities.

- **Native cloud simplicity.** Historically, security teams had to make tradeoffs when it came to cloud firewalls. Firewalls from cloud service providers offered better ease of use but could only be deployed in their own environment, which is a non-starter when most organizations support hybrid, multicloud infrastructure. Conversely, third-party firewalls deployed as VMs were more difficult to manage. Firewall platforms supporting managed firewall services can remove the operational burden of infrastructure management from security teams and enable them to focus on more strategic activities. This requires deep support and integrations across all CSPs, which was cited by 31% of organizations as a key attribute.

© 2026 TechTarget, Inc. All rights reserved. Unauthorized reproduction prohibited.

This Omdia White Paper was commissioned by Palo Alto Networks and is distributed under license from TechTarget, Inc.

- **Purpose-built security for AI workloads.** While scale, management, and simplicity are important, security cannot be sacrificed. Network firewalls today must offer protections not only for traditional cloud environments but also for containers and especially dynamic AI workloads, which both traditional firewalls and CSP firewalls were not designed to protect. Nearly half of organizations (47%) indicated security for AI models and applications was key firewall platform attribute, making it the most common response cited.

Figure 3: Key attributes of firewall platforms

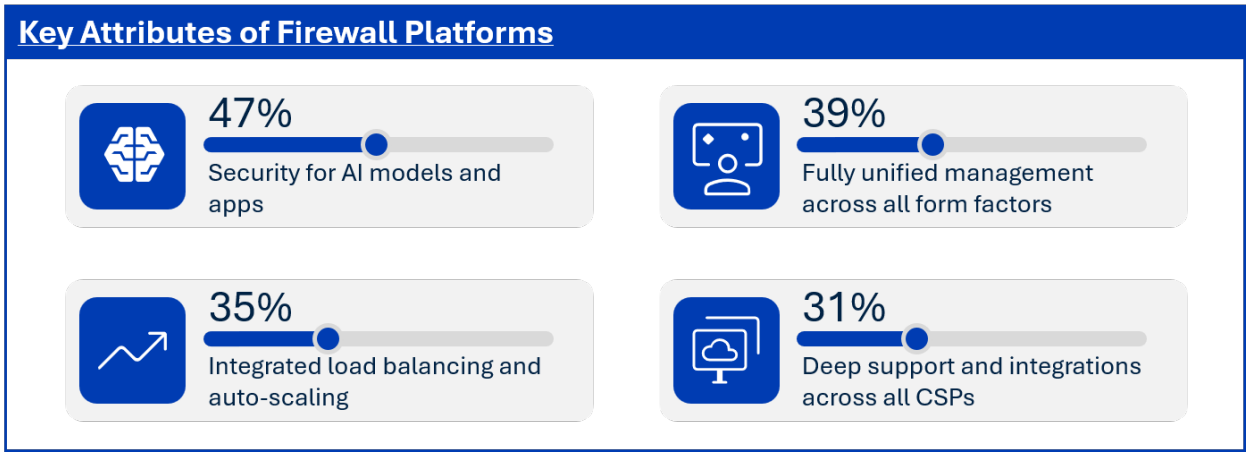

Source: Omdia

Modernizing firewall architectures with Palo Alto Networks Software Firewalls

Palo Alto Networks Software Firewall Platform helps security teams modernize their network security approach for multicloud, AI-centric environments. The platform is built around four pillars: continuous discovery, automated deployment, comprehensive protection, and operational command.

The platform helps security teams eliminate blind spots by continuously discovering every application, workload, AI service, and traffic path. Moreover, it identifies missing firewall coverage and CSP firewalls that teams may want to replace. Finally, Palo Alto Network's Cloud Network and AI Risk Assessment (CLARA) delivers actionable insights into cloud and AI risk across the environment. It includes a cloud network risk assessment, AWS and Azure firewall benchmarking, and an AI risk assessment.

© 2026 TechTarget, Inc. All rights reserved. Unauthorized reproduction prohibited.

This Omdia White Paper was commissioned by Palo Alto Networks and is distributed under license from TechTarget, Inc.

Once teams have a picture of the environment, critical gaps, and their greatest risks, they can automatically deploy the right enforcement point where protection is needed. The platform includes Cloud NGFW for managed cloud firewalling and VM-series for customer-managed virtual firewalling for protection of traditional workloads. For advanced use cases, including AI workloads and containerized environments, Palo Alto offers AI Runtime Security for self-managed deployments and Managed AI Runtime Security for Palo Alto managed cloud deployments.

The Palo Alto Networks Software Firewall platform is built around four pillars: continuous discovery, automated deployment, comprehensive protection, and operational command.

Comprehensive protection is provided through Palo Alto Network's Cloud Delivered Security Services (CDSS). This includes advanced threat prevention, DLP, advanced URL filtering, and more. Because Palo Alto Networks supports the same protection across its entire network security portfolio, including AI Runtime Security, customers enjoy consistent security across all workloads, including large-scale AI workloads, containers, and east-west traffic. This includes full Layer 7 security across north-south and east-west traffic, helping to reduce lateral movement risk and secure cloud environments at scale.

Perhaps most importantly, unified operational command helps teams manage their entire network security posture from a single command center. This offers centralized policy, logging, posture, and health across the entire environment, all managed through Strata Cloud Manager.

The core architecture of the platform is Palo Alto Network's Multicloud Security Fabric (MSF). This enables security policies across AWS, Azure, GCP, and private clouds to be standardized from a single control plane (Strata Cloud Manager), eliminating operational silos and manual network stitching. By simplifying and unifying protection across different locations throughout the environment, security teams can reduce operational complexity with centralized multicloud control, automate protection deployment without added overhead, and ensure consistent security posture across all cloud environments.

To address the scale needed for high throughput and AI traffic spikes, Palo Alto's Hyperscale Security Fabric (HSF) delivers elastic scale up to 200Gbps with active N+1 session resilience. HSF enables the platform components to automatically scale with unpredictable AI traffic bursts, eliminates bottlenecks from manual scaling and capacity limits, and removes idle capacity with intelligent load-balanced scaling.

© 2026 TechTarget, Inc. All rights reserved. Unauthorized reproduction prohibited.

This Omdia White Paper was commissioned by Palo Alto Networks and is distributed under license from TechTarget, Inc.

As noted earlier, applying network security to containers has been an ongoing struggle for many security teams. To solve for this, Palo Alto Networks has developed Container Network Security, which provides purpose-built runtime protection for Kubernetes and microservices. It integrates natively via CNI chaining to deliver deep Layer 7 inspection and inline threat prevention for pod-to-pod and cluster-to-AI traffic without having to deploy firewalls inside every cluster.

Figure 4: Palo Alto Networks unified security for multicloud and AI

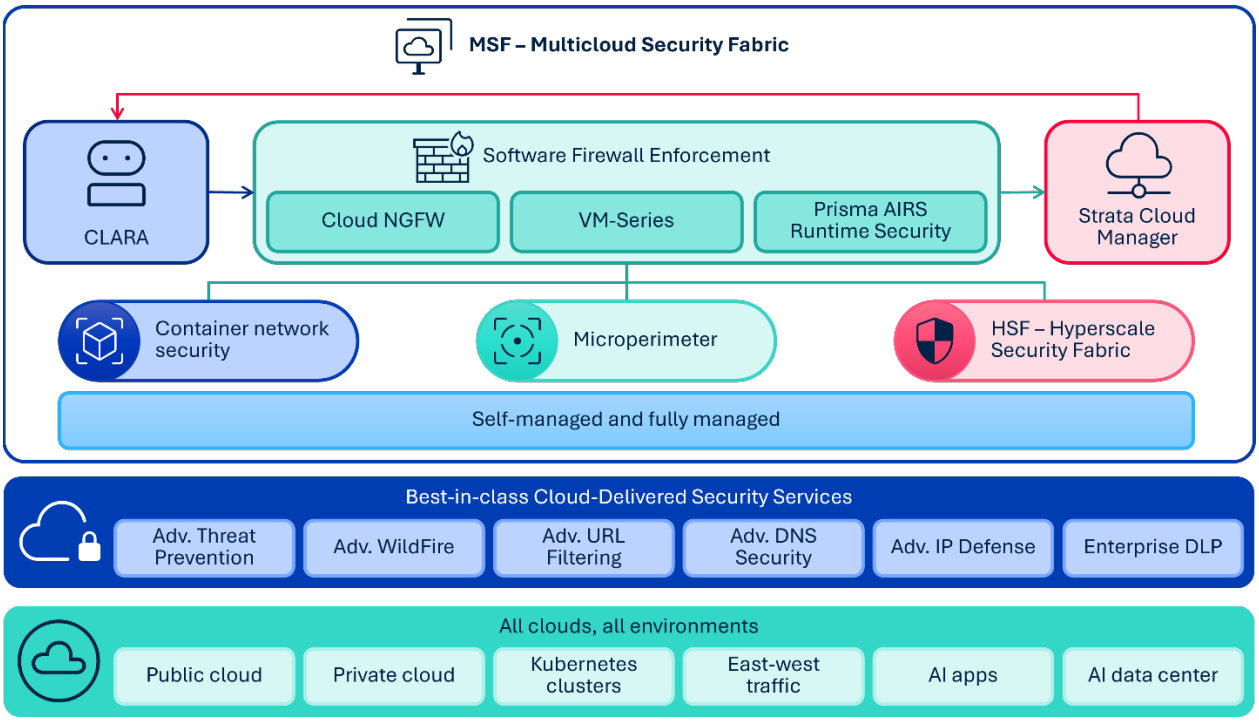

Source: Omdia

Finally, to enforce selective, Layer 7-aware microsegmentation around critical workloads, Microperimeter offers a lightweight PAN Traffic Redirector. This programmatically routes designated east-west traffic to software firewalls, containing lateral threat movement without hypervisor dependency or complex network rearchitecting. It enforces Layer 7 microperimeters around every workload, pod, and autonomous agent, providing deep inspection for east-west traffic that traditional firewalls miss, stopping lateral movement across containers and cloud environments.

While all these innovations help operationally, some organizations still want a fully managed solution. Cloud NGFW provides this as a fully managed firewall as a service (FWaaS), which integrates natively into popular cloud provider environments such as AWS,

© 2026 TechTarget, Inc. All rights reserved. Unauthorized reproduction prohibited.

This Omdia White Paper was commissioned by Palo Alto Networks and is distributed under license from TechTarget, Inc.

Azure, GCP, and more. This helps organizations that have traditionally used CSP firewalls for ease of use benefit from Palo Alto Network's enterprise-grade security efficacy, while also enjoying lower TCO compared to self-managed virtual appliances. Further, the new Managed AI Runtime Security offering extends the same zero-operational-overhead benefits to advanced firewall and AI workload protection.

By combining Multicloud and Hyperscale Fabrics with deep container network security, microperimeters, and fully managed deployment options, Palo Alto turns the network into a single control point, eliminating the compromise between AI runtime security and operational agility. Palo Alto Network's approach to securing AI infrastructure (combining purpose-built AI security capabilities with existing security infrastructure) helps prevent security teams from being the bottleneck to innovation.

© 2026 TechTarget, Inc. All rights reserved. Unauthorized reproduction prohibited.

This Omdia White Paper was commissioned by Palo Alto Networks and is distributed under license from TechTarget, Inc.

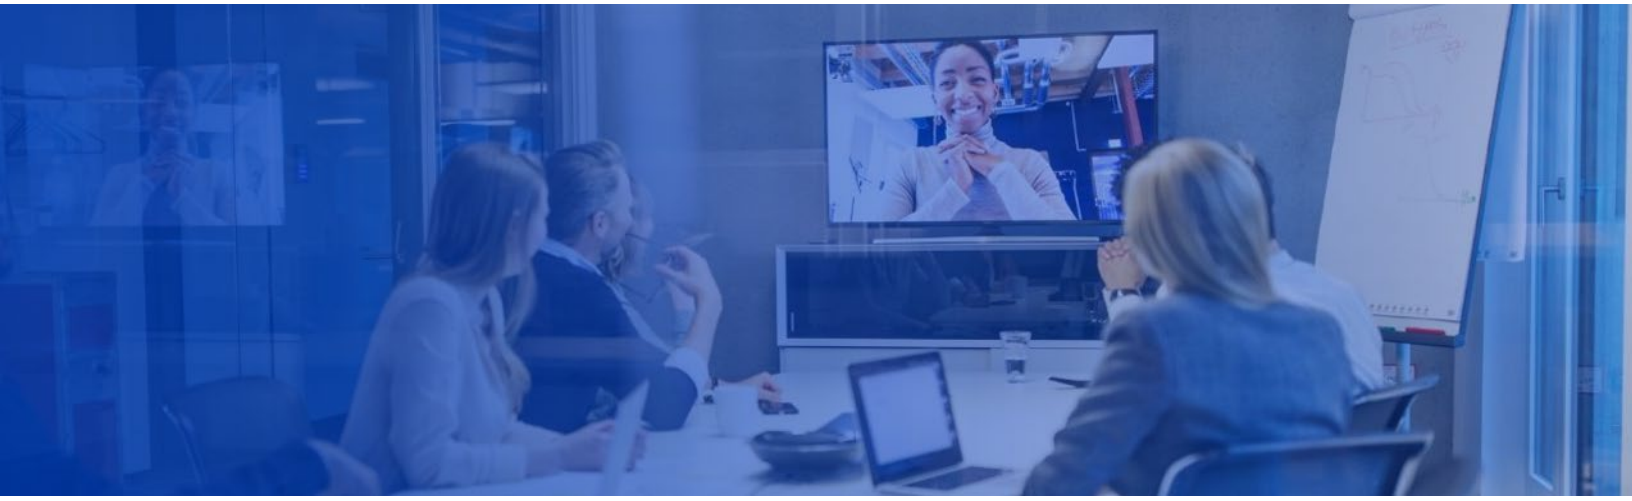

Conclusions

Security and IT leaders must begin to transition from siloed, perimeter-centric controls to a platformized, hybrid mesh firewall architecture. Modern enterprise transformation—whether driven by multicloud expansion, containerized microservices, or emerging AI initiatives—cannot succeed on an outdated, fragmented security foundation. As cloud networks become increasingly complex, security leaders must evolve their architectures beyond point tools and consolidate onto a unified software firewall platform. By deploying Multicloud and Hyperscale Security Fabrics, enforcing container network security through traffic redirection, and extending specialized protection to AI workloads via AI runtime security, enterprises can transform their network into a single control point, protecting every cloud, workload, and microperimeter while accelerating digital innovation safely.

This Omdia White Paper was commissioned by Palo Alto Networks and is distributed under license from TechTarget, Inc.

Omdia consulting

Omdia is a market-leading data, research, and consulting business focused on helping digital service providers, technology companies, and enterprise decision makers thrive in the connected digital economy. Through our global base of analysts, we offer expert analysis and strategic insight across the IT, telecoms, and media industries.

We create business advantage for our customers by providing actionable insight to support business planning, product development, and go-to-market initiatives.

Our unique combination of authoritative data, market analysis, and vertical industry expertise is designed to empower decision-making, helping our clients profit from new technologies and capitalize on evolving business models.

Omdia is part of Informa TechTarget, a B2B Materials information services business serving the technology, media, and telecoms sector. The Informa group is listed on the London Stock Exchange.

We hope that this analysis will help you make informed and imaginative business decisions. If you have further requirements, Omdia's consulting team may be able to help your company identify future trends and opportunities.

Get in touch

www.omdia.com
askananalyst@omdia.com

© 2026 TechTarget, Inc. All rights reserved. Unauthorized reproduction prohibited.

Copyright notice and disclaimer

The Omdia research, data, and information referenced herein (the "Omdia Materials") are the copyrighted property of TechTarget, Inc. and its subsidiaries or affiliates (together "Informa TechTarget") or its third-party data providers and represent data, research, opinions, or viewpoints published by Informa TechTarget and are not representations of fact.

The Omdia Materials reflect information and opinions from the original publication date and not from the date of this document. The information and opinions expressed in the Omdia Materials are subject to change without notice, and Informa TechTarget does not have any duty or responsibility to update the Omdia Materials or this publication as a result.

Omdia Materials are delivered on an "as-is" and "as-available" basis. No representation or warranty, express or implied, is made as to the fairness, accuracy, completeness, or correctness of the information, opinions, and conclusions contained in Omdia Materials.

To the maximum extent permitted by law, Informa TechTarget and its affiliates, officers, directors, employees, agents, and third-party data providers disclaim any liability (including, without limitation, any liability arising from fault or negligence) as to the accuracy or completeness or use of the Omdia. Informa TechTarget will not, under any circumstance whatsoever, be liable for any trading, investment, commercial, or other decisions based on or made in reliance of the Omdia Materials.