Futurum RESEARCH

HPE



# Converged by Necessity: Why the Network Is the Control Point for Zero Trust in the AI Era

NOVEMBER 2026



# Executive Summary

**The rapid evolution of frontier and open-weight AI models on cybersecurity tasks is giving attackers growing capability to exploit flaws faster than patch cycles can respond.**

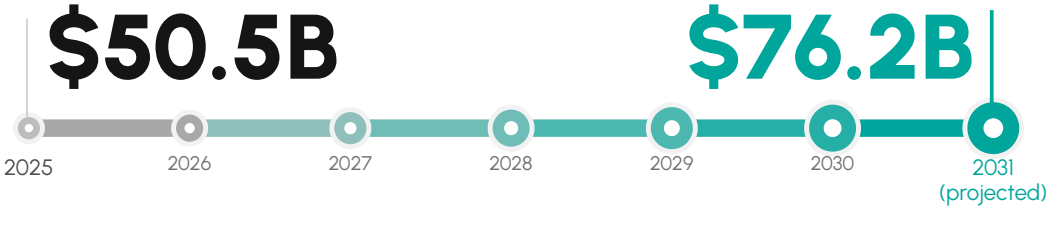
That raises the stakes for the 80% of organizations already fielding significant incidents, and they are still investing: 67% expect their cybersecurity budgets to grow over the next year.[1] Access to the most capable models is often restricted, but that restriction tends to wane, and prudent organizations should plan for that capability to reach attackers more broadly. When finding and exploiting a flaw can be largely automated, traditional security programs built on practices such as prioritized patching must evolve. Patching remains necessary, but it can no longer be the primary defense.

As AI accelerates the discovery and exploitation of vulnerabilities, an important point is often overlooked: organizations can't control every threat or vulnerability, but they can control their exposure by limiting what attackers can reach across the network. That makes the network the natural enforcement point for managing exposure, often under the zero trust umbrella. Secure access service edge (SASE) is key here, essentially where convergence starts. Zero trust holds only when it runs end-to-end, from campus access points and switches through the WAN to the cloud, and reaches the IoT and unmanaged devices that cannot run an agent.

Most vendors now describe their offerings as converged. This report sets out nine questions that test whether convergence is real at the policy layer. It closes with an overview of HPE Networking's approach.



## Network security market




**7.1%** annual growth rate

**#1** largest product segment in cybersecurity

**15%** of total cybersecurity spend

1 Futurum Research. (2026). 1H2026 Cybersecurity Decision-Makers Survey. N=929 security decision-makers.

# Motivated Attackers Are Gaining the Ability to Outpace Patching

Finding a flaw, building an exploit, and testing it against a target have always taken skilled people time. Frontier AI is compressing all three. Access to the most capable models is often restricted, but that restriction tends to wane over time, and open-weight models, which are far more accessible, are not far behind in capability.

The time it takes to find and exploit a flaw has long been one of the dampening effects on cyber risk: the gap between a flaw's discovery and its exploitation gives patch programs room to work. When discovery, weaponization, and testing can be automated and run in parallel, that gap narrows sharply. A patch cycle measured in weeks struggles to keep pace with an adversary able to act in hours.

None of this means every attack is imminent. The shift is that motivated attackers can increasingly move faster, and security programs should plan for the ones who do.

The pressure lands on programs already stretched. In Futurum's survey of security decision-makers, 80% reported at least one significant security incident in the past 12 months, and 18% reported six or more. Budget growth, mostly in the 5% to 15% range, will not by itself buy back the time that faster attackers take away.[1]

Exposure is one of the variables defenders still control: whether a vulnerable system can be discovered and reached at all. An application endpoint that is not reachable does not respond, so it cannot be scanned, tested, or exploited from outside, whatever flaws it carries. Reducing reachability is a key preventive measure, and one that keeps its value as attackers get faster.

## Two levers
## Where defenders can still gain ground



### Patch speed

Compressing as attacker capability grows

**Manual discovery and exploit building**

Time to exploit: weeks

Patch cycle fits inside the window

**AI-automated discovery, weaponization, testing**

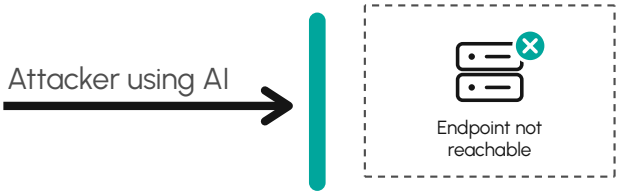
Time to exploit: hours

**Patch cycle no longer fits: the system stays exposed**



### Exposure reduction

Limits attacks on AI-discovered vulnerabilities




If a system cannot be discovered or reached, it does not respond, so it cannot be scanned, tested, or exploited from outside, whatever flaws it carries.

# Convergence Has to Start at the Network

If exposure is one of the variables defenders control, the network is where an organization acts on it. Whether a system answers, who can find it, and what path leads to it are decided by how traffic is routed, inspected, and permitted. Security tools beside the network can observe exposure; zero trust enforcement built into the network can shrink it at the source.

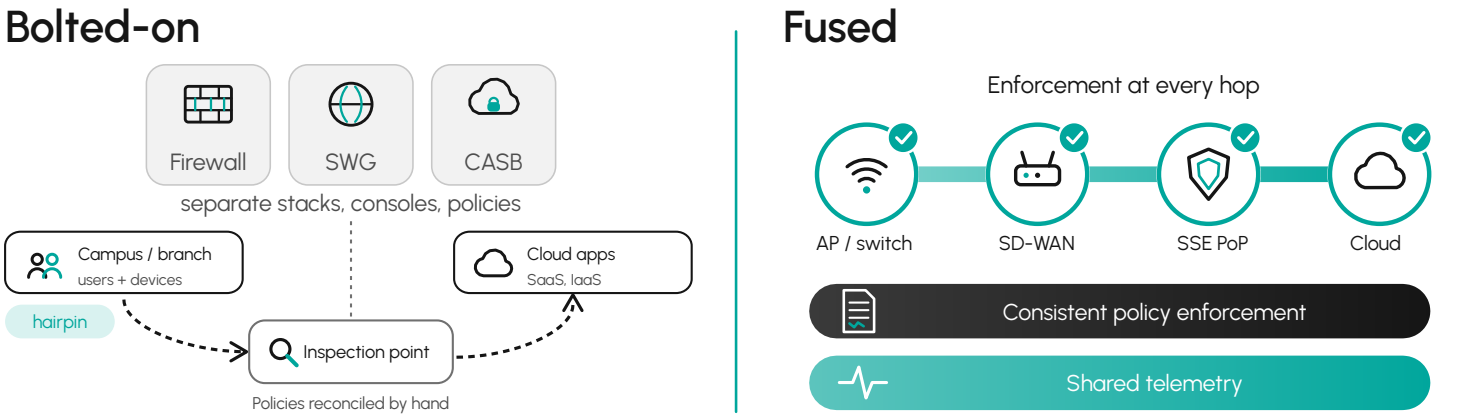
Buyers are spending accordingly. Futurum projects that network security, already the largest product segment in cybersecurity at 15% of total spend, will grow to $76.2 billion by 2031, a 7.1% annual rate.[2] In Futurum's network security module, 62% of decision-makers expect their network security budget to grow over the next 12 months.[3]

SASE has emerged as a leading architectural approach for acting on exposure. It combines software-defined WAN (SD-WAN) connectivity with cloud-delivered security capabilities, including zero trust network access (ZTNA), secure web gateway (SWG), and cloud access security broker (CASB), under one policy model. Futurum defines the category as the convergence of wide-area networking and security functions to securely connect users, systems, endpoints, and branch offices to applications and data, regardless of location. Success is measured by consistent user experience, inline threat prevention, and the operational ease of enforcing policy at scale.[4]

SASE is deeply tied to zero trust, the principle that access is minimized to what each user and device needs and that traffic is continuously evaluated. Zero trust delivers the most when it runs end-to-end, from the access points and switches in the branch and campus, across the WAN, and into the data center and cloud. That means managing the full stack (wired, wireless, and WAN) as one environment with one set of policies, with AI-driven operations that correlate events across layers. Disconnects between network and security teams, extending consistent policy to remote sites, and implementing zero trust without disrupting productivity rank closely together among buyers' top challenges.[3]

**Architecture**

## Bolted-on vs. fused enforcement




Timing favors acting now. Many organizations are rebuilding their networks to carry AI workloads, and building enforcement into that rebuild usually costs less than adding it afterward.

Convergence carries a trade-off. Consolidating policy into one plane concentrates risk: a single control plane becomes a high-value target, and its availability becomes a security dependency. Buyers should weigh that deliberately.

2 Futurum Research. (2026, May). 1H2026 Cybersecurity Market Sizing and Five-Year Forecast. Base case.
3 Futurum Research. (2026). 1H2026 Cybersecurity Decision-Makers Survey, network security module. N=103 security decision-makers (SASE friction: N=100 evaluating or migrating to SASE). Respondents ranked their top three.
4 Futurum Research. (2026, March 26). Futurum Signal Report: Secure Access Service Edge (SASE). Document FSRSASE202603.

# Zero Trust Stops Short of the Devices That Cannot Run an Agent

SASE as first defined protects the branch, the remote user, and the cloud application. It does much less for what already sits inside the campus. Printers, cameras, badge readers, building controls, medical and industrial equipment, contractor laptops, and personal devices share the network with managed endpoints, and most cannot run an endpoint agent. SASE access that depends on an agent leaves a large and hard-to-govern share of the device population outside the policy.

Closing that gap means applying zero trust at the point of connection. Network access control (NAC) applies the same identity-and-posture logic to every device that connects, managed or not: each is identified, profiled, and segmented with only the access its role requires. The approach that holds up at scale has three properties:



**Continuous, AI-Driven Profiling:**
Machine learning identifies device types and watches behavior over time, so risk posture is reassessed as behavior changes.



**Enforcement Built Into the Infrastructure:**
Access points, switches, and SD-WAN edge devices apply policy themselves, with no overlay that traffic must be steered toward.



**One Policy Language Across Local and Remote Access:**
A contractor's laptop gets the same decision on campus or from home.

Devices that cannot run an agent still reach the internet, so web protection has to extend to them through a dedicated gateway tunnel from the network edge that requires nothing on the device.

Practitioners describe this work in operational terms. Customer feedback on a widely deployed NAC platform returns to the same jobs: keeping unknown devices off the network, segmenting IoT, and applying one policy across wired and wireless.

> *"[With NAC], we have been able to streamline the number of SSIDs that are used in our district from about 7 down to 3 main networks for staff and students, IoT and guests."*
>
> — Network access control user, US school district (G2 verified review)[5]

5 G2. (2026). Verified user reviews, HPE Networking ClearPass Policy Manager (listed on G2 as "HPE Aruba Networking SSE"). Accessed September 2026.

Others describe a platform that "allows you to profile any device (domain, IoT, etc.) specifically and assign networks dynamically to it," and the end of manual work: "No more logging into the network switch to configure VLANs on a specific port."[5] More than half of decision-makers (54%) in Futurum's network security module already report NAC in widespread or localized use.[3]

## Coverage
# What each layer reaches

### SASE (agent-based)

Protects users and sites that can run an agent

- Remote users
- Branch offices
- Managed laptops
- Cloud apps

Leaves out: devices that cannot run an agent

### Infrastructure-embedded NAC

Every connecting device, identified, profiled, and segmented at the point of connection

- Managed laptops
- IoT: cameras, printers, badge readers
- Operational technology (OT)
- Building controls
- Medical & industrial equipment
- Contractor laptops
- BYOD
- Guests

Together: zero trust end to end, from campus access points and switches through the WAN to the cloud

Campus access points — Switches — WAN — Cloud



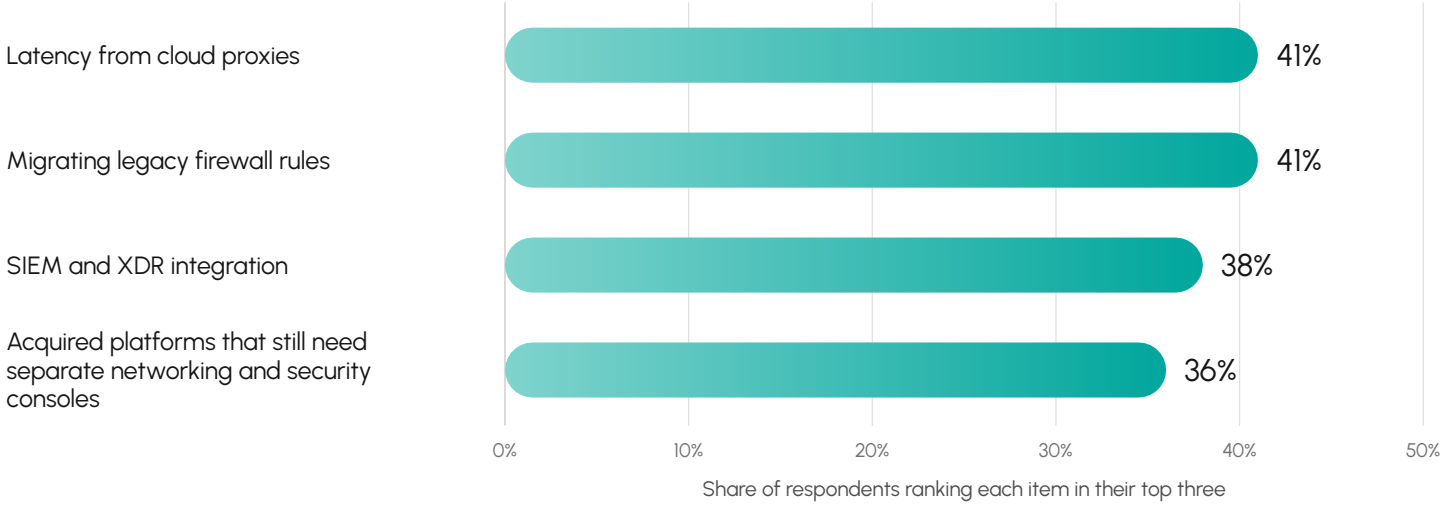


# Nine Questions Separate Real Convergence From Superficial Integration

Converged is now the default claim in this market. These questions evaluate whether a SASE offering delivers genuine architectural and operational convergence at the policy layer, and they apply to any vendor in the category. They also track what buyers report slowing them down (see Figure 1). Among decision-makers evaluating or migrating to SASE, the most-cited friction points are latency from cloud proxies (41%), migrating legacy firewall rules (41%), SIEM and XDR integration (38%), and platforms assembled through acquisitions that still need separate networking and security consoles (36%).[3]

**Figure 1**

## What Slows SASE Consolidation




Note. N=100 security decision-makers evaluating or migrating to SASE. Respondents ranked their top three friction points; the four most-cited are shown, and percentages do not sum to 100%. Data from 1H2026 Cybersecurity Decision-Makers Survey, network security module, by Futurum Research (2026).[3]

Most of these friction points show up during migration and day-to-day operations, when changing course is costly. The nine questions below bring them into the evaluation instead. Each targets a design choice that is hard to reverse after deployment: how policy is modeled, where traffic is inspected, which devices are covered, how the platform is run, and how it behaves when something fails. Use them with every vendor on the shortlist, and ask to see each answer working in a demonstration or a reference deployment.

1. **Is there one SASE console?** A unified console over separate networking and security management simplifies the view and policy management. Futurum's SASE Signal research finds that portfolios marketed as SASE but requiring separate consoles and fragmented policy enforcement face growing market rejection.[4] Ask how and where policy is enforced across various domains (wired, wireless, WAN, and cloud security services).

2. **Is traffic inspected once, on the best path?** Running each service from its own point of presence (PoP), or decrypting TLS separately for each, adds latency at every step. A single-pass engine applying ZTNA, web, and SaaS controls together keeps performance predictable across users, applications, and AI workloads. Ask whether SD-WAN path selection is coordinated with security inspection, so traffic reaches the right services by application, identity, device posture, and risk, without manual configuration or unnecessary backhauling.

3. **How does zero trust extend to devices that cannot run an agent?** Ask how unmanaged and IoT devices are identified and segmented, and whether they share a policy language with managed devices and remote access.

4. **When do the management planes merge?** Where networking and security are still managed separately, ask for a dated roadmap to bring them together.

5. **What does it take to operate?** Skills gaps rank among buyers' leading network security challenges (28% top-three).[3] Ask about day-one deployment, migrating existing firewall rules, and keeping policy consistent across distributed sites.

6. **What happens when the policy plane is unreachable?** Centralized enforcement makes availability a security property, because an unreachable policy plane can hold up access across the environment. Ask about failure modes and local survivability for policy enforcement.

7. **Does the platform feed the SOC?** Enforcement telemetry should flow into existing SIEM and XDR workflows without custom integration.

8. **How does the platform ensure local traffic processing and security enforcement?** Ask if it supports local security enforcement, minimizes traffic backhauling to the cloud, and addresses data residency and regulatory requirements.

9. **What is the AI in the console allowed to do?** An assistant that explains a problem differs from an agent that identifies the root cause and remediates on its own. Ask how broadly it operates (wired, wireless, and WAN), what it can change, what it logs, and who approves it.

Reducing reachability is a strong control, and it has limits. It does not stop an attacker using stolen credentials, a compromised device already inside the trust boundary, or abuse of SaaS-to-SaaS connections, and the access broker itself becomes a target worth defending. Convergence also raises switching costs. It works best as the foundation of a layered program, with identity, endpoint, and detection controls built on top.



# HPE Networking Builds SASE Up From the Network

HPE Networking builds its SASE approach from the network outward, the model this report describes (see Table 1).

**SD-WAN Foundation:** HPE Networking EdgeConnect has evolved from an SD-WAN offering into a SASE platform that brings SD-WAN and Security Service Edge (SSE) together in a single console. Futurum's SASE Signal research credits its SD-WAN capabilities with strong traffic steering and bandwidth use over broadband connections, and with business value from reducing reliance on MPLS.[4] HPE cites business intent overlays, tunnel bonding, and path conditioning among the features behind that performance, and traffic steering extends across multicloud environments.

**SD-WAN and SSE Integration:** HPE says the connectivity between EdgeConnect and HPE Networking SSE is automated end-to-end, with no manual tunnel configuration, policy translation, or synchronization, which reduces deployment risk and time.[6] AppExpress monitors application experience and shifts each application to the best-performing path, including traffic bound for SSE, while Smart Routing connects users to the best-performing PoP.[6]

**Consistent Policy Enforcement:** HPE says SD-WAN and SSE policies are centrally managed through the EdgeConnect interface, with a single policy engine spanning SSE capabilities such as SWG, ZTNA, and CASB. Across the network, shared user identities, device roles, and contextual intelligence help enable consistent security policy enforcement.

**Inspection and Private Access:** A private-access connector embedded in EdgeConnect reaches private applications over the SD-WAN fabric, with no separate connector infrastructure, and HPE Networking Private Edge provides in-boundary access without hairpinning through cloud PoPs. HPE says all SSE services run on a single policy engine within a single PoP, so traffic is decrypted once and inspected in a single pass. Inline data loss prevention (DLP) works across SWG and CASB to keep sensitive data from leaving through web and SaaS traffic.[6]

**Branch Firewall:** HPE says EdgeConnect's built-in firewall, with intrusion detection and prevention (IDS/IPS), machine learning-based adaptive distributed denial-of-service (DDoS) protection, URL filtering, and role-based segmentation, can replace legacy branch firewalls and reduce hardware at the site.[6] Role-based segmentation also limits east-west traffic between devices, the internal side of reducing exposure.

6 HPE. (2026, June 16). HPE Networking EdgeConnect unifies SD-WAN and SSE in an AI-native console to power SASE. HPE Community.

## Table 1
# HPE Networking SASE Capabilities at a Glance

| Capability area | What it does |
|---|---|
| SD-WAN foundation | EdgeConnect brings SD-WAN and SSE together in a single console; business intent overlays, tunnel bonding, and path conditioning; traffic steering across multicloud environments |
| SD-WAN and SSE integration | Automated end-to-end connectivity with no manual tunnel configuration, policy translation, or synchronization; AppExpress shifts each application to the best-performing path; Smart Routing connects users to the best-performing PoP |
| Consistent policy enforcement | Single policy engine for SWG, ZTNA, and CASB, managed centrally in EdgeConnect; shared user identities, device roles, and context across the network |
| Inspection and private access | All SSE services in one policy engine within a single PoP, so traffic is decrypted once and inspected in one pass; embedded private-access connector; HPE Networking Private Edge for in-boundary access; inline DLP across SWG and CASB |
| Branch firewall | Built-in firewall with IDS/IPS, machine learning-based adaptive DDoS protection, URL filtering, and role-based segmentation that limits east-west traffic |
| AI in operations | SASE Copilot answers natural-language questions and speeds troubleshooting; Marvis AI extends across wired, wireless, and WAN through HPE Mist and HPE Aruba Central integration |
| Universal zero trust | ClearPass Policy Manager, Central NAC, and Mist Access Assurance apply AI-driven device profiling and microsegmentation through access points, switches, and SD-WAN; dedicated SWG tunnel for agentless devices |
| Zero-day defense | SRX Firewalls in a hybrid mesh: inline, proxyless AI-predictive threat prevention and machine learning-generated signatures |

Note. Descriptions summarize HPE product materials and reflect HPE's stated capabilities. Data from HPE[6–8]; SD-WAN assessment from Futurum Signal Report: Secure Access Service Edge (SASE), by Futurum Research (2026).[4]

**AI in Operations:** The EdgeConnect console is AI-native.[6] SASE Copilot, built into EdgeConnect, answers natural-language questions, surfaces analytics, and speeds troubleshooting, which HPE positions as a way to ease the skills burden on operators.[7] For day-two operations, HPE is bringing EdgeConnect into HPE Mist and HPE Aruba Central. With the Mist and Central integrations, HPE says Marvis AI is extending AI-driven operations across wired, wireless, and WAN, adding predictive insights, recommended actions, and increasingly automated remediation.

**Universal Zero Trust:** HPE's aim is to apply zero trust principles to managed and unmanaged devices alike, with consistent identity, policy, and enforcement. HPE Networking ClearPass Policy Manager, HPE Aruba Networking Central NAC, and HPE Mist Networking Access Assurance apply AI-driven device profiling and microsegmentation enforced through access points, switches, and SD-WAN, which HPE calls universal ZTNA. A dedicated SWG tunnel from EdgeConnect extends web protection to agentless devices. Futurum's SASE Signal research also notes the flexibility of HPE's agentless-first ZTNA model.[4]

**Zero-Day Defense:** HPE Networking SRX Firewalls extend the architecture as a hybrid mesh. HPE says their AI-predictive threat prevention inspects traffic inline through a lightweight, proxyless architecture rather than waiting for a full file download, and uses machine learning to detect anomalous behavior and generate signatures, often before vulnerabilities are formally documented.[8]

7 HPE. (2026). SASE Copilot in HPE Networking EdgeConnect. hpe.com/psnow/doc/a50016073enw
8 HPE. (2026). Modernizing security and connectivity with unified SASE (solution overview). hpe.com/psnow/doc/a00156430enw

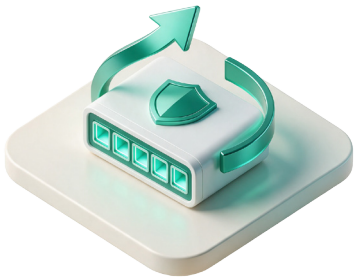


# Four Moves for the Next Planning Cycle



### Map exposure before the next purchase

Inventory what is reachable from outside and what connects inside, including devices that cannot run an agent. That map sets the scope for every other decision.



### Put enforcement into the network refresh

Write zero trust requirements into the next network design and RFP, so security is designed into access, WAN, and cloud connectivity from the start.



### Run every vendor through the nine questions

Weight one policy model and agentless coverage most heavily, because those separate architectural convergence from a shared console.



### Treat convergence as the foundation

Keep identity, endpoint, and detection controls layered on top, and plan for resilience if any single control fails.

Attacker speed is outside any organization's control. What an attacker can reach is largely a network design decision, and it should be made deliberately.

# Information About This Report

## AUTHORS

**Fernando Montenegro**
VP & Practice Lead, Cybersecurity & Resilience | The Futurum Group

## PUBLISHER

**Futurum Research**

## INQUIRIES

Contact us if you would like to discuss this report and The Futurum Group will respond promptly.

## CITATIONS

This paper can be cited by accredited press and analysts, but must be cited in context, displaying author's name, author's title, and "The Futurum Group." Non-press and non-analysts must receive prior written permission by The Futurum Group for any citations.

## LICENSING

This document, including any supporting materials, is owned by The Futurum Group. This publication may not be reproduced, distributed, or shared in any form without the prior written permission of The Futurum Group.

## DISCLOSURES

The Futurum Group provides research, analysis, advising, and consulting to many high-tech companies, including those mentioned in this paper. No employees at the firm hold any equity positions with any companies cited in this document.

HPE

## ABOUT HPE

HPE Networking builds its SASE approach from the network outward, unifying SD-WAN and Security Service Edge (SSE) capabilities into a single AI-native console. By integrating EdgeConnect with HPE Networking SSE, the solution delivers automated end-to-end connectivity without requiring manual tunnel configuration or policy translation. Traffic is continuously monitored and routed along optimal paths while undergoing single-pass inspection for ZTNA, SWG, and CASB policies. Additionally, HPE extends universal universal zero trust principles across both managed and unmanaged devices, using AI-driven profiling and microsegmentation. This comprehensive model enables consistent policy enforcement and enhanced threat prevention from branch locations to the cloud.

Futurum®

## ABOUT THE FUTURUM GROUP

The Futurum Group is an independent research, analysis, and advisory firm, focused on digital innovation and market-disrupting technologies and trends. Every day our analysts, researchers, and advisors help business leaders from around the world anticipate tectonic shifts in their industries and leverage disruptive innovation to either gain or maintain a competitive advantage in their markets.

Futurum®
**CONTACT INFORMATION:** The Futurum Group LLC | futurumgroup.com | (833) 722-5337

© 2026 The Futurum Group. All rights reserved.