

Zero Trust Platforms

Alexei Balaganski


July 13, 2026



LEADERSHIP
COMPASS
2026

Leadership Compass

Zero Trust Platforms



This KuppingerCole Analysts Leadership Compass provides an overview of the Zero Trust Platforms market and a compass to help you find the solution that best meets your needs. It examines vendors delivering integrated platforms that unify access control, segmentation, and contextual enforcement across users, workloads, devices, services, and AI-driven systems in hybrid enterprise environments. It provides an assessment of how well these platforms operationalize Zero Trust principles of explicit verification, least privilege, and continuous trust evaluation across on-premises, cloud, SaaS, and edge infrastructures.

Contents

Introduction	4
Key Findings.....	5
Market Analysis.....	6
Delivery Models	7
Required Capabilities	8
Leadership	10
Overall Leadership	10
Product Leadership	13
Innovation Leadership.....	14
Market Leadership.....	15
Product and Vendor Evaluation.....	17
AppGate – AppGate ZTNA	19
archTIS – NC Protect	23
Broadcom – Symantec SSE.....	27
Cisco – Cisco Secure Access	31
Fortra – Fortra Cloud Security Platform.....	35
GoodAccess – GoodAccess ZTNA Platform	39
Illumio – Illumio Breach Containment Platform.....	43
Ivanti – Ivanti Neurons for Zero Trust Access	47
Microsoft – Microsoft Security	51
NetFoundry – NetFoundry.....	55
Netskope – Netskope One Platform	59
Pomerium – Pomerium	63
Sophos – Sophos Workspace Protection (ZTNA)	67
Specops – Specops Device Trust.....	71
Systancia – cyberelements	75

Teleport – Teleport Infrastructure Identity Platform	79
ThreatLocker – ThreatLocker Zero Trust Platform	83
Unosecur – Unosecur	87
Whiteswan – Whiteswan Identity Security Platform	91
Zero Networks – Zero Networks Platform.....	95
Vendors to Watch	99
1Password.....	99
Absolute Security	99
Akamai Technologies.....	100
Aryaka.....	100
Barracuda Networks.....	101
Cato Networks.....	101
Citrix Systems (Cloud Software Group).....	102
Cloudflare	102
Ergon Informatik (Airlock).....	102
Forum Systems	103
Google	103
Jamf.....	104
Palo Alto Networks.....	104
Sangfor Technologies	104
Skyhigh Security.....	105
Tailscale	105
Twingate.....	106
Versa Networks	106
Zscaler.....	107
Related Research	108
Copyright.....	108

Introduction

Zero Trust has evolved from an ambitious security principle into one of the defining architectural paradigms for modern enterprise security. What began as a reaction against the obsolete assumption of a trusted internal network has become a broad strategic framework for controlling access, reducing lateral movement, and continuously evaluating trust across highly distributed digital environments.

The traditional enterprise perimeter has effectively disappeared. Organizations now operate across combinations of on-premises infrastructure, public cloud platforms, SaaS ecosystems, edge deployments, and partner-operated environments. Employees, contractors, applications, APIs, workloads, and increasingly AI-driven systems interact continuously across infrastructure boundaries that security teams neither own nor fully control. In such environments, implicit trust based on network location is no longer defensible.

In response, many organizations have adopted individual Zero Trust technologies over the past several years. Multifactor Authentication (MFA), Zero Trust Network Access (ZTNA), adaptive authentication, microsegmentation, endpoint posture validation, and cloud workload controls have all become important elements of modern security programs. However, these technologies are often implemented independently through disconnected control planes.

This fragmented approach exposes the limitations of tool-centric Zero Trust strategies. Security policies become inconsistent across environments, risk signals are not shared effectively between systems, and operational complexity increases as organizations accumulate overlapping tools and integrations. Attackers who compromise legitimate credentials, tokens, or machine identities can often still move laterally across hybrid environments with limited resistance.

Zero Trust Platforms (ZTPs) have emerged to address this architectural gap. Zero Trust itself remains a security strategy built around explicit verification, least privilege, and continuous evaluation of trust. ZTPs do not replace this strategy. Instead, they provide the technical and operational foundation required to implement Zero Trust principles consistently and at scale.

These platforms unify policy orchestration, contextual analysis, segmentation, and enforcement across users, workloads, devices, services, APIs, and increasingly non-human and AI-driven identities. Rather than treating identity, network security, workload protection, and access management as separate domains, Zero Trust Platforms create a shared policy and enforcement fabric spanning hybrid infrastructures.

This shift represents an important evolution in enterprise security architecture. Early Zero Trust initiatives focused heavily on securing remote user access through ZTNA technologies that replaced traditional VPNs. While this was a critical first step, modern attack paths increasingly target east-west traffic, service-to-service communication, API abuse, cloud-native workloads, and machine identities. Zero Trust, therefore, must evolve beyond user

authentication and network access into a universal enforcement model for all digital interactions.

The market is evolving rapidly. Some vendors approach Zero Trust Platforms from a networking and segmentation heritage, expanding ZTNA into broader identity-aware connectivity and microsegmentation capabilities. Others originate from IAM, cloud security, SASE, workload protection, or data security backgrounds and are building integrated enforcement fabrics around those strengths.

In this Leadership Compass, we assess vendors delivering integrated Zero Trust Platform capabilities designed to operationalize Zero Trust architectures across modern enterprise infrastructures. We examine their strengths, innovation, deployment flexibility, interoperability, and strategic direction in a market that is rapidly becoming foundational to enterprise cybersecurity strategy. To better understand the fundamental principles this report is based on, please refer to [KuppingerCole Analysts' Research Methodology](#).

Key Findings

- Zero Trust Platforms are emerging as the operational foundation for applying Zero Trust consistently across hybrid environments, addressing inconsistent enforcement, operational complexity, and visibility gaps.
- Identity is now the primary enterprise attack surface, with attackers routinely abusing legitimate credentials, tokens, APIs, and machine identities.
- ZTNA alone is no longer sufficient: buyers increasingly require unified enforcement across users, workloads, applications, and east-west traffic.
- Continuous, adaptive authorization is replacing static authentication. 75% of vendors covered in this report continuously re-evaluate active sessions, and 70% trigger re-evaluation on identity risk signals.
- Standard-based federation is now a baseline capability (85% SAML, 80% OIDC), but phishing-resistant authentication remains a differentiator: only 60% natively support passkeys despite NIST and CISA prioritization.
- Just-in-time access has moved from the privileged-access niche into mainstream Zero Trust (80% support JIT or time-bound access), and 85% support graduated enforcement instead of binary allow/deny decisions.
- Device posture signals are well integrated (80% MDM/UEM, 75% EDR/XDR), but cloud posture is not: 40% still cannot evaluate CSPM signals when making access decisions.
- Non-human identity is the next frontier. 80% of vendors integrate with secrets or workload identity systems, but only 40% natively issue workload identity certificates, leaving significant headroom for SPIFFE/SPIRE architectures.
- Generative AI is shifting from a marketing claim to a product feature: 65% use it to improve usability and 55% to author or refine policies.
- SaaS is the dominant delivery model (80%), and most vendors offer both agent-based and agentless options; cryptographic tenant isolation, increasingly required by regulated industries, is available in fewer than half (45%).

- SIEM is the universal interoperability factor (90%), and 80% of platforms integrate with SIEM or SOAR for automated response actions, including device containment.
- The market is converging across IAM, ZTNA, microsegmentation, cloud security, API security, and data-centric security, with non-human identities now outnumbering human users in many enterprises.
- Buyers should evaluate Zero Trust Platforms on architectural integration, policy consistency, deployment flexibility, interoperability, and machine-identity support rather than on isolated feature sets.

Market Analysis

Over the past several years, Zero Trust evolved from an aspirational security philosophy into a widely accepted strategic objective. Most enterprises today claim to be pursuing some form of Zero Trust initiative. However, many implementations remain fragmented collections of point products rather than cohesive operational architectures. This fragmentation is one of the primary drivers behind the emergence of ZTPs.

Organizations now operate across highly distributed hybrid environments. Traditional network-centric security models cannot provide sufficient visibility or control across these environments. At the same time, the operational burden created by disconnected identity, endpoint, network, and cloud security tools has become increasingly difficult to manage.

As a result, enterprises are shifting from isolated Zero Trust controls toward integrated platforms that unify policy orchestration, contextual evaluation, segmentation, and enforcement. Several major trends are accelerating this transition.

First, identity has become the new primary attack surface. Credential theft, token replay, session hijacking, API abuse, and supply-chain compromise campaigns increasingly rely on abusing legitimate access rather than exploiting vulnerabilities. Attackers no longer need to “break in” when they can authenticate themselves using compromised identities or services.

Second, the adoption of cloud-native architectures has dramatically increased east-west traffic and service-to-service communication. Applications now consist of distributed microservices communicating dynamically across cloud platforms and APIs. Traditional perimeter segmentation models are poorly suited for these environments. Identity-aware segmentation and continuous authorization that operate independently of network topology are required to address the challenge.

Third, automation and non-human identities are expanding rapidly. Workloads, APIs, service accounts, Kubernetes clusters, CI/CD pipelines, IoT devices, and AI agents now represent a large and growing percentage of enterprise identities. In many environments, machine identities outnumber human users by orders of magnitude. Zero Trust can no longer focus primarily on authenticating employees connecting remotely to applications. It must also

govern how workloads communicate with each other, how APIs exchange sensitive data, and how autonomous systems access resources.

Another major trend shaping the market is the convergence of historically separate security domains. IAM, PAM, ZTNA, microsegmentation, cloud workload protection, SASE, SSE, API security, and data-centric security have traditionally evolved as distinct markets with separate operational teams and management models. ZTPs increasingly act as the connective layer between these domains.

Leading vendors are building architectures that integrate identity signals, device posture, network telemetry, workload context, behavioral analytics, and policy orchestration into unified enforcement models. Some emphasize identity-centric architectures, while others focus on network segmentation, workload-level controls, or cloud-native enforcement. However, the overall direction of the market is clearly toward policy convergence and centralized contextual enforcement.

The market remains highly diverse in terms of vendor positioning. Large established cybersecurity vendors increasingly market integrated Zero Trust architectures spanning identity, networking, cloud security, and endpoint protection. Specialized vendors continue to innovate aggressively in areas such as identity-aware microsegmentation, workload-level policy enforcement, AI-driven contextual analytics, and machine identity protection.

AI is beginning to influence this market in multiple ways. Vendors increasingly use ML for anomaly detection and behavioral analytics and Generative AI (GenAI) for policy recommendations, adaptive access decisions, and automated threat response. GenAI systems and autonomous agents are creating entirely new security requirements around authorization, segmentation, and machine-initiated access.

Operational complexity remains one of the largest barriers to successful Zero Trust adoption. Many organizations still struggle with fragmented ownership between networking, IAM, cloud, infrastructure, and security teams. Legacy applications and infrastructure continue to complicate deployment, particularly in large hybrid environments.

Finally, the market itself is still evolving conceptually. The Zero Trust narrative is no longer just about replacing VPNs or enforcing MFA. It is becoming the foundational operational model for securing distributed digital environments. ZTPs represent the technical implementation of that shift.

Delivery Models

ZTPs are predominantly delivered as cloud-native SaaS offerings, reflecting the distributed nature of modern enterprise environments and the need for globally scalable policy enforcement. SaaS delivery enables vendors to provide continuous updates, distributed enforcement points, centralized policy orchestration, and rapid deployment across hybrid infrastructures.

However, deployment flexibility remains a critical requirement for many organizations. Enterprises operating in regulated industries or highly sensitive environments often require private cloud, customer-managed Infrastructure as a Service (IaaS), hybrid, or fully on-premises deployment models to address compliance, sovereignty, operational, or latency requirements.

Thus, many vendors support combinations of SaaS, containerized, virtual appliance, Kubernetes-native, and software-based deployments. Some platforms rely heavily on cloud-delivered control planes with distributed enforcement nodes, while others offer self-managed architectures designed for isolated or highly customized environments.

Architectural approaches also vary significantly across the market. Some solutions use proxy-based or brokered access models commonly associated with ZTNA and SSE architectures. Others rely on agent-based segmentation, embedded workload enforcement, API-centric controls, or identity-aware gateways. Increasingly, platforms combine multiple enforcement methods depending on use case and deployment constraints.

Licensing models differ widely as well. Vendors may charge per user, per workload, per application, per protected resource, per transaction volume, or according to network throughput and infrastructure scale. Organizations evaluating ZTPs should carefully assess how pricing models align with long-term expansion plans, machine identity growth, API usage, and multi-cloud adoption.

Required Capabilities

The term “Zero Trust Platform” itself is used inconsistently across the industry. Many vendors present ZTNA, SSE, or SASE offerings as ZTPs despite offering only limited contextual enforcement or lacking unified policy orchestration. For this reason, organizations must look beyond marketing terminology and focus on architectural depth.

The following is a list of criteria for evaluation in this Leadership Compass on Zero Trust Platforms:

- Secure, identity-aware connectivity for users, workloads, services, APIs, and devices
- Continuous verification and adaptive trust evaluation based on identity, context, posture, and behavioral signals
- Fine-grained access control and authorization policies supporting least privilege and just-in-time access
- Identity-aware segmentation and microsegmentation for limiting lateral movement
- Centralized policy orchestration with distributed enforcement across hybrid environments
- Support for strong authentication methods including MFA and passwordless authentication, and FIDO2, WebAuthn, OIDC, and federation standards
- Integration with IAM, PAM, endpoint security, cloud security, SIEM, SOAR, and DevOps ecosystems
- Device, workload, and application posture evaluation using contextual telemetry

- Real-time monitoring, analytics, logging, and anomaly detection across sessions and workloads
- Protection for machine identities, including workloads, APIs, service accounts, containers, and AI agents
- Support for both human and Non-Human Identity (NHI) governance and access enforcement
- Threat prevention capabilities, including detection of credential abuse, anomalous access, and lateral movement
- APIs, SDKs, and policy-as-code interfaces for automation and developer integration
- Support for cloud-native, hybrid, private cloud, and on-premises deployments
- Scalability for globally distributed enterprises and latency-sensitive workloads
- Encryption and protection of data in transit and at rest
- Comprehensive audit logging and reporting for regulatory compliance
- Open integration capabilities and interoperability across heterogeneous enterprise environments
- Administrative dashboards, workflow orchestration, and operational analytics
- High availability, resiliency, and distributed enforcement architecture

The following items are considered additional and innovative capabilities:

- AI-driven behavioral analytics and adaptive policy recommendations
- Continuous authorization and event-driven access reevaluation during active sessions
- Semantic or intent-aware controls for AI agents and automated systems
- Unified policy models spanning users, workloads, APIs, and data access
- Deep workload identity protection and service-to-service authorization
- Runtime protection for cloud-native workloads and Kubernetes environments
- Risk-adaptive access enforcement using behavioral, environmental, and threat intelligence signals
- Advanced identity-aware microsegmentation independent of network topology
- Support for policy-as-code and DevSecOps integration pipelines
- AI-assisted policy creation, analytics, and investigation workflows
- Data-centric policy enforcement integrated directly into Zero Trust decisions
- Fine-grained controls for east-west traffic across hybrid and multi-cloud infrastructures

Leadership

Selecting a vendor of a product or service must not be based only on the information provided in this KuppingerCole Analysts Leadership Compass. The Leadership Compass provides a comparison based on standardized criteria and can help identify vendors that should be considered for further evaluation. However, a thorough selection process should include a detailed analysis and a Proof of Concept (PoC) or pilot phase, based on the specific criteria of the customer.

Overall Leadership

The Overall Leadership chart reflects how vendors balance product strength, innovation, and market presence. The chart evaluates vendors in this market across two dimensions. The horizontal axis measures product leadership, and the vertical axis measures innovation leadership.

Overall leaders are those in the upper right-hand quadrant, with the bubbles colored red. Vendors in this quadrant score highly in both dimensions, earning overall leadership status. These organizations ship mature products while maintaining strong innovation.


The bubble size reflects each vendor's relative strength in terms of market leadership. Further details on market leadership can be found in the Market Leadership section.

Product leaders are found to the right of the vertical dividing line. The vendors in the lower right quadrant, colored black, are established vendors with full-featured products but fewer innovative differentiators. These providers demonstrate solid execution using conventional methods. Further details on product leadership can be found in the Product Leadership section.

Innovation leaders are found above the horizontal dividing line. Vendors in the upper left-hand quadrant, colored black, have advanced technology approaches but less mature product offerings. These companies often introduce new technical approaches that are not yet widely adopted. Further details on innovation leadership can be found in the Innovation Leadership section.

Vendors in the lower left-hand quadrant, colored grey, are challengers for both product and innovation leadership in this market. This quadrant is not uniform. Because bubble size reflects market presence, it includes some established vendors with large customer bases and broad geographic reach whose products are less differentiated in depth or innovation within this specific segment, alongside smaller, regional, or specialist vendors focused on particular use cases or geographies. Challenger positioning on these two axes is not a statement about a vendor's overall scale, financial strength, or relevance for specific requirements.

The Overall Leaders in the Leadership Compass on Zero Trust Platforms are Microsoft, Teleport, Cisco, Broadcom, Netskope, NetFoundry, Appgate, Zero Networks, and Illumio.



Microsoft is recognized as an overall leader because it combines broad product coverage, strong innovation, and unmatched market presence across identity, endpoint, cloud, data protection, SIEM/SOAR, and AI governance. Its strength lies in the depth of integration across its entire portfolio, making it especially relevant for organizations already standardized on Microsoft cloud and security platforms.

Teleport demonstrates overall leadership through its strong infrastructure identity architecture, short-lived certificate model, mature access workflows, and depth of coverage. Its focus on secretless access and machine identity governance aligns closely with the

direction of the Zero Trust Platforms market, even though it complements rather than replaces broader SSE or SASE platforms.

Cisco combines broad product depth with substantial market reach and a strong innovation trajectory. Cisco Secure Access, Duo, Secure Workload, Hypershield, and Splunk together form a layered Zero Trust architecture spanning user access, identity intelligence, endpoint posture, workload segmentation, threat analytics, and AI access controls.

Broadcom's Symantec SSE is positioned as an overall leader because of its mature SSE architecture, strong ZTNA capabilities, embedded DLP, threat prevention, and integration with Symantec Data Center Security for workload segmentation. Its use of Google's private backbone, mature policy engine, and strong compliance posture make it a comprehensive enterprise offering.

Netskope demonstrates overall leadership through the breadth of the Netskope One platform, combining SSE, SD-WAN, ZTNA, Cloud Access Security Broker (CASB), Secure Web Gateway (SWG), DLP, Data Security Posture Management (DSPM), UEBA, AI security, Remote Browser Isolation (RBI), Enterprise Browser, and device intelligence under a unified architecture. Its private security cloud, data protection depth, and AI governance capabilities make it especially relevant for organizations looking for cloud-native, data-aware Zero Trust enforcement.

NetFoundry is recognized for its workload, API, OT/IoT, and machine-to-machine Zero Trust architecture. Its OpenZiti-based overlay, endpoint-generated private keys, service-level microsegmentation, SDK embedding model, and MCP and LLM gateway capabilities provide a strong answer to modern non-human identity and agentic AI access challenges.


AppGate is recognized because of its mature direct-routed ZTNA architecture, strong policy engine, deployment flexibility, and proven suitability for regulated, government, OT, and air-gapped environments. Its per-user micro-firewall model, strong automation interfaces, and operational resilience make it a platform for organizations that prioritize control, sovereignty, and high-assurance access.

Zero Networks earns the overall leader recognition through its highly automated approach to microsegmentation and privileged network access control. Its agentless model, automated policy generation, privileged port and service account controls, and RPC firewall capabilities address some of the hardest lateral movement problems in real enterprise environments.

Illumio is an overall leader because of its depth in Zero Trust segmentation and breach containment. Its label-based policy model, host-native enforcement, broad workload coverage, process-level segmentation, and strong compliance posture make it a leading platform for reducing lateral movement across hybrid and multi-cloud environments.

Product Leadership

Product leadership is the first specific category examined below. This view is based on the presence and completeness of required features as defined in the required capabilities section above. The chart is horizontal and divided into two areas: the vendors to the right of the chart and colored red are product leaders, and those to the left and colored black are product leadership challengers.



Product Leadership in the Zero Trust Platforms market reflects the completeness, maturity, and integration depth of required capabilities across secure connectivity, access policy, context and posture evaluation, segmentation, monitoring, threat prevention, data-centric controls, and AI and machine identity protection. The strongest products are not just collections of tools: they provide coherent policy models, distributed enforcement, strong interoperability, and automated operational workflows.

The Product Leaders are Microsoft, Teleport, Cisco, Netskope, Broadcom, Fortra, AppGate, NetFoundry, Illumio, and Zero Networks.

Microsoft leads this group with broad functional depth across identity, device posture, cloud security, data governance, SIEM/SOAR, workload identity, and AI agent governance. Cisco, Netskope, and Broadcom follow closely, each combining mature SSE or access enforcement with broader ecosystem integration, threat prevention, and data or workload security capabilities.

Teleport stands out for infrastructure identity, secretless access, and machine identity governance, while Fortra is particularly strong in data-centric security through DLP, DRM, classification, and DSPM. AppGate, NetFoundry, Illumio, and Zero Networks provide strong product depth in more specialized but strategically important areas: direct-routed ZTNA, workload and API connectivity, breach containment, and automated microsegmentation.


The remaining vendors are positioned as Challengers in Product Leadership. Many of them have strong capabilities in specific areas, such as device posture, data-centric enforcement, endpoint control, European sovereignty, or identity risk visibility, but do not yet provide the

same breadth, maturity, or architectural integration as the leaders. We recommend referencing the vendor-specific chapters to understand how these offerings may still be relevant for particular use cases or deployment environments.

Innovation Leadership

Innovation Leadership in this market is shaped by how effectively vendors address the next generation of Zero Trust requirements. These include continuous and event-driven authorization, machine and workload identity, AI agent governance, policy-as-code, identity-aware microsegmentation, risk-adaptive enforcement, data-centric policy controls, and operational use of AI for policy management, investigation, and response.

The innovation chart is horizontal and divided into two areas: the vendors to the right of the chart and colored red are innovation leaders, and those to the left and colored black are innovation leadership challengers.



The Innovation Leaders are Microsoft, Teleport, Cisco, Broadcom, Netskope, NetFoundry, Zero Networks, Pomerium, Illumio, and AppGate.

Microsoft is recognized for its work on security and governance controls for agentic AI across its entire portfolio. Teleport demonstrates innovation through its certificate-centric architecture, workload identity, and access governance for human and non-human infrastructure identities. Cisco, Broadcom, and Netskope all show strong momentum in AI access controls, data-aware policy, and broader SSE convergence.


NetFoundry and Pomerium stand out for developer-oriented Zero Trust enforcement models that extend naturally to APIs, workloads, and AI agent interactions. Zero Networks is innovative in applying MFA directly at the OS firewall layer for privileged ports and in automating microsegmentation through deterministic policy generation.

Illumio continues to innovate in breach containment through AI-enhanced security graph analytics and workload-centric microsegmentation, while AppGate remains distinctive through its direct-routed architecture and automation-driven entitlement model.

Other vendors are positioned as Challengers in Innovation Leadership. Several show meaningful innovation in narrower domains, including data-centric access control, device trust, endpoint containment, European sovereign access platforms, and identity risk analytics.

Market Leadership

Market Leadership reflects customer base, geographic reach, partner ecosystem, deployment scale, financial strength, and the ability to support global enterprise adoption. In this category, scale matters more than technical differentiation alone. A vendor may be highly innovative or technically strong while still remaining a challenger in market leadership if its reach, installed base, or partner ecosystem is more limited.



The Market Leaders are Microsoft, Cisco, Broadcom, Fortra, Sophos, and Ivanti.

Microsoft, Cisco, and Broadcom lead Market Leadership because of their sheer enterprise footprints, global partner networks, and relevance for large, regulated organizations. Fortra is positioned as a leader due to its established customer base and presence across data protection, secure access, and managed security channels.

Sophos maintains strong market leadership through its global channel and MSP ecosystem, particularly across SMB and mid-market segments. Ivanti's market leadership is supported by its installed base in endpoint management and secure access, as well as its practical migration path from legacy VPN to ZTNA.

Challengers show credible market traction and growing visibility but remain behind the leaders in scale, reach, ecosystem maturity, or overall market influence. They are substantial market participants, but their leadership position is still consolidating.

Contenders are stable vendors with recognizable market presence in specific regions, customer segments, or use cases. They may have loyal customers and sustainable business models but generally lack the global reach or strategic visibility of market leaders.

Emerging vendors are companies with promising growth trajectories but still limited market penetration. They may be gaining attention quickly, but their customer base, channel maturity, brand recognition, and enterprise references are still developing.

Entrants are early-stage or narrowly positioned vendors that are beginning to establish market relevance. Their position is typically limited by smaller customer bases, narrower geographic coverage, lower visibility, or the need to demonstrate repeatable growth at scale.

Product and Vendor Evaluation

This section provides a structured evaluation of each product and vendor included in this KuppingerCole Analysts Leadership Compass. Each profile contains a description of the company and its offering, an analysis of relevant capabilities for this market segment, and an assessment of strengths and challenges. Where applicable, vendors are positioned as leaders in the product, innovation, or market categories.

In addition to these standard Leadership Compass categories, we provide detailed capability ratings for every vendor. These ratings are visualized in a spider chart that reflects performance across the functional and technical criteria defined for this market segment. The spider chart complements the written analysis and enables direct comparison across solutions.

For this Leadership Compass, the capability categories evaluated are described below. The first four categories are standard across all Leadership Compasses, while the latter eight are specific to the Zero Trust Platforms market segment.

Security: This covers the degree of security provided by the product or service. For products, it examines the robustness of authentication, authorization, encryption, policy enforcement, and protection mechanisms against unauthorized access, credential abuse, lateral movement, and compromise of identities or workloads. It also considers the vendor's security practices, certifications, vulnerability management, and the ability to support resilient Zero Trust architectures across hybrid environments.

Deployment: This covers how easily the solution can be deployed, operated, scaled, and maintained across enterprise environments. It evaluates support for SaaS, hybrid, private cloud, containerized, and on-premises deployment models, as well as architectural flexibility, operational complexity, scalability, and the ability to integrate into existing infrastructures without requiring disruptive redesigns.

Interoperability: This covers the ability of the solution to integrate with heterogeneous enterprise ecosystems, including IAM, PAM, endpoint security, SIEM, SOAR, cloud security, DevOps, and networking platforms. It considers support for open standards, APIs, SDKs, federation protocols, and policy interoperability across cloud, on-premises, and multi-vendor environments.

Usability: This covers the quality and consistency of both administrative and end-user experiences. It evaluates the intuitiveness of management consoles, policy authoring interfaces, dashboards, workflows, reporting, and operational controls, as well as the ability to simplify Zero Trust operations without creating unnecessary friction for users, administrators, or developers.

Secure Connectivity Enforcement: This category examines the platform's ability to provide secure, identity-aware connectivity across users, workloads, devices, applications, and services. It includes encrypted communications, authenticate-before-connect models,

elimination of implicit network trust, distributed enforcement, and secure access controls spanning hybrid and multi-cloud environments.

Access Management Policy: This measures the platform's ability to define, orchestrate, and enforce granular access policies based on identities, attributes, risk, context, and behavioral signals. It includes support for adaptive authorization, least privilege, just-in-time access, continuous policy re-evaluation, and centralized policy management with distributed enforcement.

Context and Posture: This category evaluates how effectively the solution incorporates contextual information into trust decisions. It includes user, device, workload, application, network, and environmental posture assessment, as well as integration with telemetry sources such as EDR, XDR, MDM/UEM, and cloud security platforms to support adaptive access controls and continuous trust evaluation.

Segmentation and Lateral Movement: This category examines the ability to isolate workloads, applications, services, and network flows to reduce attack surfaces and prevent unauthorized east-west movement. It considers identity-aware microsegmentation, workload-level isolation, dynamic policy enforcement, and segmentation models that operate independently of static network topology.

Monitoring and Analytics: This measures the visibility, telemetry collection, reporting, and analytical capabilities of the platform. It includes monitoring of sessions, workloads, traffic flows, policy decisions, and anomalous activity, as well as dashboards, behavioral analytics, threat detection, investigation support, and integration with external security analytics and response platforms.


Threat Prevention: This category evaluates the platform's ability to detect, block, and mitigate threats targeting identities, workloads, applications, APIs, and services. It includes protections against credential abuse, unauthorized access, malware propagation, session hijacking, API misuse, insider threats, and suspicious lateral movement, as well as integration with broader detection and response ecosystems.

Data-Centric Security: This category examines the platform's ability to protect sensitive data through policy-driven controls integrated into Zero Trust enforcement. It includes encryption, data access governance, inspection of data flows, redaction, DLP integration, and context-aware restrictions on how sensitive information is accessed, transmitted, or processed across distributed environments.

AI and Machine Identity Protection: This category measures how effectively the platform secures non-human identities, including workloads, APIs, service accounts, containers, devices, and AI agents. It includes machine identity governance, service-to-service authorization, workload authentication, least-privilege enforcement for automated systems, and controls designed to manage AI-driven or machine-initiated access and actions.

AppGate – AppGate ZTNA

AppGate



Founded in 2000 and headquartered in Coral Gables, Florida, AppGate is an enterprise cybersecurity company specializing in Zero Trust Network Access. Its eponymous flagship product is complemented by AppGate Risk Sentinel for dynamic risk-based access control and Application Discovery for AI-powered policy recommendations derived from observed access patterns. The platform has a well-established presence in regulated industries, government, and defense environments, with support for demanding operational and compliance requirements.

AppGate ZTNA is built on a direct-routed architecture, which differentiates it from the many cloud-proxied ZTNA services in the market. Rather than routing traffic through a vendor-operated cloud broker, the platform establishes individual Mutual TLS (mTLS)-encrypted

tunnels directly between the client and the gateway at the enforcement point. Protected resources and AppGate infrastructure are cloaked using Single Packet Authorization (SPA), which drops unauthenticated connection attempts before network exposure occurs. The architecture separates the control plane, handled by the Controller cluster as the policy decision point, from the data plane, which remains in the customer environment whether controllers are self-hosted or managed by the vendor. Tunnels default to TLS 1.3, with Datagram Transport Layer Security (DTLS) and Quick UDP Internet Connections (QUIC) available as alternatives.

The platform supports multiple enforcement models. Agent-based clients are available for managed endpoints across Windows, macOS, Linux, ChromeOS, iOS, and Android. A browser-based clientless portal supports unmanaged and BYOD access, while a headless client addresses server and workload use cases. AppGate ZTNA also provides a Kubernetes sidecar injector and connector appliances for IoT, OT, and legacy network segments where agent installation is not feasible. Support for air-gapped deployments and configurable tolerance for WAN disruption makes the platform relevant for restricted, remote, and operationally demanding environments.

Access policy combines role-based and attribute-based controls evaluated at session establishment using user and device claims. The Controller issues cryptographically signed entitlement tokens to clients, which gateways translate into per-user micro-firewall rules. In effect, every user receives a dynamically constructed “segment of one,” with access limited to resources for which valid entitlements exist. Policy conditions that can be defined include time-based restrictions, step-up MFA, and runtime re-evaluation triggered by claim changes, risk alerts, location or network changes, and behavioral anomalies. Just-in-time (JIT) access is supported through integrations with IT Service Management (ITSM) platforms, where access to sensitive resources can depend on an open, assigned ticket and be revoked automatically when the ticket is closed. Approval workflows are handled externally through these integrations. Policy authoring is available through the administrative UI, a REST API, a Terraform provider, and JavaScript-based scripting for advanced logic.

Context and posture evaluation draws on device claims collected by the client, including running processes, file presence, registry entries, and custom script outputs. Through Risk Sentinel, AppGate integrates with EDR and MDM platforms such as CrowdStrike, SentinelOne, and Microsoft Intune, feeding risk scores into the policy engine and enabling mid-session re-evaluation or access revocation. Cloud resolver integrations allow policies to reference cloud metadata and tags dynamically, so access to workloads can adjust as infrastructure is provisioned, repurposed, or terminated without manual policy updates. Posture requirements can be defined independently for each entitlement, allowing different resources to have different contextual conditions enforced.

Segmentation is identity-driven and independent of network topology. Entitlements target logical hostnames and resolver objects rather than IP ranges, and gateways enforce default-deny access by installing only explicit allow rules per user session. Workload-to-workload and service-to-service segmentation are supported through headless clients, connector appliances, and the Kubernetes sidecar injector. Application discovery observes real user access patterns and generates policy recommendations to help organizations move from

initially broad access rules toward fine-grained least privilege. AppGate ZTNA also integrates with third-party microsegmentation tools, including Illumio and ColorTokens, for environments that need stronger workload isolation.

Monitoring and analytics cover session-level audit logs, per-user firewall state, policy decision records, and appliance health metrics exposed through Prometheus and SNMP. Real-time session visibility is available in the administrative console, including user claims, device claims, and active entitlements for troubleshooting and incident investigation. Telemetry can be streamed externally with documented integrations for Splunk, Azure Monitor, Datadog, Elasticsearch, OpenSearch, and others. AppGate ZTNA does not perform native behavioral baselining or anomaly detection. Automated responses include session termination, token revocation, step-up authentication, and account or device blocking.

Threat prevention is primarily architectural. SPA, mutual TLS, short-lived signed tokens, and per-user micro-firewall rules reduce the exposed attack surface and constrain lateral movement at the session level. Port scan detection is included. The platform does not natively provide deep packet inspection, malware scanning, or API schema validation, relying on third-party solutions that customers have to procure separately for these functions. AppGate's Risk Sentinel takes in signals from EDR platforms and, upon evaluation, can dynamically reduce privileges or force re-authentication when endpoint threat indicators change during a session.

The same identity and context-driven access model applies to NHIs. Service accounts, Kubernetes workloads, and CI/CD agents are supported through the headless client and Kubernetes sidecar injector, receiving the same entitlement-based micro-firewall enforcement as human users. Machine identities are issued with controller-signed certificates and tokens that can be rotated or revoked immediately. Least-privilege enforcement for automated systems uses the same policy engine. However, AppGate ZTNA does not currently provide intent-level controls over agent actions, so policies determine which resources a machine identity can reach but do not constrain what it does inside a session. AI workload isolation is supported through network-level segmentation and access control, with graph-based threat detection for AI agents added to the platform recently.

AppGate holds SOC 2 Type II attestation covering security, availability, processing integrity, confidentiality, and privacy, as well as FIPS 140-3 compliance. Deployment flexibility is a genuine strength: the platform supports fully self-managed, AppGate-hosted, and hybrid models; runs on physical hardware, virtual machines, or cloud infrastructure across AWS, Azure, GCP, and Oracle Cloud; and operates in air-gapped and limited-connectivity environments. Support languages are limited to English, Spanish, and Japanese, and documentation is available only in English. A set of REST APIs powers the entire management UI and enables full programmatic control, including operation without the admin console. The end-user experience is designed to be largely transparent after initial setup.

The company's roadmap focuses on AI-driven operations, including Observability AI for end-to-end user experience monitoring and access issue diagnosis and Optimizer AI for continuous policy and entitlement analysis.

AppGate ZTNA is a technically distinctive platform best suited to organizations that prioritize performance, architectural control, and deployment flexibility over broad security service consolidation. Its direct-routed model, strong policy engine, and deployment resilience make it particularly strong in demanding enterprise, government, and critical infrastructure environments, especially those with hybrid or isolated infrastructure, OT and IoT requirements, or strict data sovereignty constraints. AppGate ZTNA is most relevant where fine-grained access control, operational resilience, and ecosystem interoperability matter more than platform bundling.


Strengths

- Direct-routed architecture eliminates inbound network exposure and cloud-proxy dependency, giving customers full traffic sovereignty
- Per-user micro-firewall model with cryptographically signed, individually constructed entitlements per session
- One of the few ZTNA solutions validated for fully air-gapped, limited-connectivity, OT, and government environments
- Application Discovery uses ML to analyze real access patterns and recommend least-privilege policy migrations
- A rich automation surface enables deep DevOps integration
- Strong Latin America presence with dedicated partner and support coverage, uncommon among specialist ZTNA vendors

Challenges

- No native data-centric security capabilities
- Behavioral threat analytics depends entirely on external SIEM and EDR integrations
- East-west workload-to-workload segmentation is not natively enforced and requires third-party tools
- Inline threat inspection is not supported
- Limited market presence outside the Americas, with a minimal customer base in APAC

archTIS – NC Protect



Founded in 2006 and headquartered in Canberra, Australia, with a US presence in Tampa, Florida, archTIS is a publicly listed provider of data-centric security software. Its portfolio has expanded through acquisition, including the addition of Spirion for sensitive data discovery and classification. The broader product portfolio includes the Spirion suite, NC Protect for dynamic data protection in Microsoft 365 and file-sharing environments, archTIS Trusted Data Integration (TDI) for policy-governed access to structured data sources, and Kojensi for classified information sharing across government environments. The product evaluated here is NC Protect, archTIS' primary Zero Trust offering. The company is a member of the Microsoft Intelligent Security Alliance and has MISA-approved integrations with Microsoft Purview Information Protection and Microsoft Sentinel.

NC Protect takes a distinctly data-centric approach to Zero Trust. It does not compete with network-centric ZTNA or connectivity platforms. Instead, it enforces Attribute-Based Access Control (ABAC) directly inside Microsoft 365 environments, including SharePoint Online, Microsoft Teams, and Exchange, as well as SharePoint Server on-premises and Windows File Shares. Rather than controlling how users connect, NC Protect controls what users can see and do with data. Authentication is delegated to Microsoft Entra ID, and the product has no independent connectivity enforcement capability. Network cloaking, mutual

authentication, and encrypted transit are inherited from the cloud infrastructure in which the platform runs. NC Protect components operate inside private Azure virtual networks with no public IP exposure, using TLS 1.3 for traffic and Azure Managed Identities for service-to-service authentication.

The core policy model is aligned with the NIST 800-162 architecture. Access decisions are made in real time by evaluating user attributes from Entra ID, such as department, clearance level, nationality, and group membership, against data attributes including sensitivity labels, content classifications, authorship, and metadata. The Policy Enforcement Point (PEP) is embedded directly in SharePoint, intercepting access requests and enforcing decisions at the individual file level. As a result, two users viewing the same document library may see different files, and even when both can access a file, their permitted actions may differ.

A user with viewing rights but accessing from an untrusted location may be limited to a read-only Secure Reader with a dynamic watermark, while the same user in a trusted office location can open and edit the document natively. Dynamic access control is supported through Microsoft Entra Continuous Access Evaluation, enabling near real-time token revocation when risk signals change. JIT and time-bound access is available through Microsoft Entra Privileged Identity Management.

Context and posture evaluation currently relies on signals from Entra ID and the Microsoft ecosystem, including device compliance through Intune, network location, identity risk from Entra ID Protection, and time of day. Direct integrations with third-party EDR or MDM platforms beyond the Microsoft stack are on the roadmap but not yet supported. The open attribute model gives the platform flexibility within its Microsoft-native scope: any attribute accessible through Entra ID tokens or Microsoft Graph can be incorporated into policy decisions.

Segmentation in NC Protect is identity-based and infrastructure-level rather than network-based. Within Azure, containers run in isolated private VNets with default-deny Network Security Group rules, while pod-to-pod communication is controlled through Kubernetes Network Policies at L3/L4. Workload-to-workload and service-to-service segmentation is supported within the Azure deployment. East-west movement monitoring is provided by Microsoft Defender for Containers, which maps detections to the MITRE ATT&CK framework, while Microsoft Sentinel provides behavioral anomaly detection through UEBA.

Monitoring and analytics are built around NC Protect's own audit logging of data access events. The platform captures every access attempt, policy decision, watermark application, and file action at the item level. Logs can be streamed to Microsoft Sentinel or Splunk for analysis and automated response. This fine-grained visibility into sensitive data access is a particular strength, especially for organizations that need forensic detail on who accessed which content, when, and under what conditions. Behavioral baselining and anomaly detection are handled by Sentinel rather than natively inside NC Protect.

Threat prevention is largely inherited from the Microsoft security stack, including Defender for Containers, Entra ID Protection, and Sentinel-driven playbooks for automated response. NC Protect's contribution is the granular audit stream it feeds into that ecosystem, providing

visibility that goes beyond native Microsoft 365 logging. Token theft and session hijacking are mitigated through Continuous Access Evaluation and Entra ID Protection.

The platform enforces file-level controls, including dynamic encryption through Microsoft Purview Information Protection or the optional NC Encrypt module with BYOK support. It provides secure read-only viewing through Secure Reader, prevention of download, copy, print, and screenshot actions, and dynamic watermarks for document and image files. Policies can reference NC Protect classifications as well as third-party labels from Microsoft Purview, Spirion, and ISEC7 Classify. Dynamic masking and redaction at the item level are on the near-term roadmap.

NC Protect can apply ABAC policies to API access based on the identity of registered applications in Entra ID, allowing administrators to restrict which automated systems can access sensitive content. However, broader NHI controls, including workload identity federation, service account lifecycle management, and comprehensive auditability of machine actions, are still on the roadmap. A near-term priority is extending policy enforcement to Microsoft Copilot, controlling what it can access and generate from protected content.

NC Protect holds FIPS 140-3 validation. It does not require SOC 2, ISO 27001, or FedRAMP certifications, as it operates within the customer's Azure tenant rather than as a hosted service. Deployment is restricted to Microsoft Azure, including commercial and government cloud configurations, but other cloud platforms are not supported.

The end-user experience is transparent, with NC Protect embedded into SharePoint and Teams without requiring client installation. The administrator interface includes a visual policy builder with conflict detection. REST APIs and Terraform are supported for deployment automation. External integrations are currently mostly unidirectional, with signals flowing into the platform but automated actions out to non-Microsoft third-party systems not natively supported. Support is available in English only.

The roadmap focuses on extending ABAC coverage across Microsoft Teams channels and chat, Microsoft Exchange, and Microsoft Copilot, with particular emphasis on controlling how agentic AI systems interact with sensitive content. As Copilot and similar tools make data access more indirect and less visible to end users, fine-grained policy enforcement at the data layer becomes increasingly important.

NC Protect is a specialized solution best suited to organizations deeply invested in Microsoft 365 that need granular, data-centric access control beyond native Microsoft capabilities. It is particularly relevant for government, defense, financial services, and regulated industries where need-to-know enforcement, classification, and compliance are primary concerns. Its value is strongest as a complementary control layer in a larger security architecture, especially where Microsoft 365 data protection remains both critical and underserved.


Strengths

- Native Microsoft 365 ABAC policy enforcement embedded directly into the platform
- Exceptional file-level enforcement granularity
- Persistent user-specific watermarking for documents and emails for forensic traceability
- Broad Secure Reader capabilities entirely within the browser without endpoint agents
- BYOK encryption support for organizations requiring customer-managed keys over protected content
- Microsoft Copilot policy enforcement addresses the gap of agentic AI accessing sensitive M365 content

Challenges

- The platform is scoped to the Microsoft ecosystem, including M365, SharePoint Server, and Windows File Shares; non-Microsoft repositories are not covered
- Third-party EDR and MDM integrations are on the roadmap but not yet available
- Comprehensive NHI governance is still on the roadmap; current NHI support is limited to ABAC over Entra ID-registered applications
- Automated response actions to non-Microsoft third-party systems are not natively supported
- Small relative to platform peers despite strong recent revenue growth and access to public capital markets

Broadcom – Symantec SSE



Founded in 1991 and headquartered in San Jose, California, Broadcom is a publicly listed global technology company spanning semiconductors, infrastructure software, and enterprise security. Its Enterprise Security Group operates under the Symantec brand and combines capabilities assembled through several acquisitions like Symantec, Luminate, and VMware. The product evaluated here is Symantec SSE, a comprehensive Security Service Edge platform combining ZTNA, Cloud SWG, CASB, DLP, Remote Browser Isolation, Cloud Firewall Service, and endpoint protection.

Symantec SSE is delivered over Google's private cloud backbone rather than relying only on public internet routing. Cloud traffic is routed through Google's fiber infrastructure to a broad global Point of Presence (PoP) footprint. For branch and data center connectivity, Express Connect uses Google Cloud Network Connectivity Center to reduce dependence on VPN

tunnels and support high-bandwidth peering. The ZTNA component uses a Software-Defined Perimeter (SDP) architecture with protected applications cloaked from the internet. Connectors deployed as Docker containers in the customer environment establish outbound-only tunnels to the cloud service, avoiding inbound exposure. Unauthorized users cannot see protected resources, and DNS returns non-routable addresses for unconfigured targets.

Both agent-based and agentless access are supported. The agentless path supports native RDP, SSH, and TCP-based applications without a browser extension, with enforcement implemented in the cloud backend rather than on the client device. A unified Enterprise Security Agent consolidates endpoint, network, ZTNA, CASB, SWG, and DLP enforcement into a single deployable component across Windows, macOS, iOS, and Android.

Access policy supports RBAC, ABAC, contextual, and risk-based models evaluated through a central cloud policy engine. Authorization granularity extends from server and application level down to API route, HTTP method, URI path, and SSH command level. Policy conditions include identity, device posture, location, device type, and behavioral risk. Continuous session re-evaluation is enforced by default, with context changes triggering access adjustment or session termination. Just-in-time and time-bound access policies are supported, including ITSM-driven workflows. Step-up authentication can be triggered based on risk thresholds or access to sensitive application paths. A policy simulation mode allows administrators to test access decisions before enforcement. Policy-as-code is supported through declarative JSON using an OPA-style approach, and a Terraform provider is available for infrastructure automation.

Context and posture evaluation signals include device posture through the Symantec agent, third-party tools such as OPSWAT for unmanaged devices, EDR and XDR telemetry from Symantec Endpoint Security, identity risk from connected identity providers, and application-level context. Posture requirements can also differ by application type, with distinct handling for web, SSH, RDP, and TCP-based resources.

Segmentation is delivered through two layers. ZTNA enforces application-level isolation by design, reducing lateral movement through resource cloaking and least-privilege connectivity. East-west segmentation inside data centers is handled by Symantec Data Center Security, a separate product providing application control, host-based intrusion prevention, workload isolation, and application allowlisting for physical, virtual, and containerized environments. Data Center Security supports workload-to-workload and service-to-service segmentation independent of network topology that builds policy baselines from observed application behavior. Integration with VMware vDefend allows ZTNA signals to trigger policy changes in VCF environments, coordinating enforcement across layers.

Broadcom's admin experience stands out for its built-in policy-hygiene tooling. The Insights Log proactively alerts administrators to misconfigurations and surfaces applications that should be linked together in policy but are not, turning the console from a passive editor into an active guide. Behavioral anomaly detection is driven by user risk scoring. Telemetry is available for external analysis through APIs, log streaming to object storage, and direct integrations with SIEM platforms including Splunk, Google Chronicle, IBM QRadar, and

Exabeam. Automated response actions include session termination, step-up authentication, account blocking, and device quarantine, with AI-powered threat inspection in real time.

Threat prevention is a core strength of the platform. Broadcom's Global Intelligence Network aggregates threat intelligence across endpoint, email, network, and web security telemetry and feeds it into SSE enforcement points. Inline inspection includes TLS termination, malware scanning, sandboxing, protocol inspection, and API schema validation, applied consistently to agent-based and agentless sessions. Remote Browser Isolation (RBI) is embedded directly into the ZTNA, rendering application content in a secure cloud container so that source code, metadata, and objects do not reach the endpoint. This reduces exposure to web application vulnerabilities without requiring a separate WAF. Token theft and session hijacking are mitigated through continuous session validation, token binding to specific devices and TLS sessions, and behavioral analytics.

Data-centric security is delivered through Symantec DLP, a mature capability embedded into the SSE platform. DLP covers ZTNA, Cloud SWG, CASB, email, and endpoint channels under a unified policy and management console. Controls include content inspection, browser session controls such as copy-paste prevention, download and upload blocking, and print prevention through RBI. Integration with Microsoft Purview supports classification-based encryption and watermarking. Masking and redaction are not natively provided, although metadata can be supplied to external tools that perform those functions.

NHI support covers workloads, APIs, service accounts, and AI agents. Machine identities receive short-lived certificates or tokens tied to specific workloads, with automated rotation and immediate revocation through policy updates. A notable recent capability is ZTNA as an MCP gateway, addressing the lost identity problem in agentic AI architectures through Secure On-Behalf-Of token exchange. This preserves the user's least-privilege identity through the MCP layer and applies inspection to both prompts and responses. Shadow AI discovery, risk assessment across known GenAI applications, and prompt injection prevention are supported for public AI services. Private LLM access control through the ZTNA data path provides inspection and enforcement for internally hosted models.

Broadcom holds ISO 27001, SOC 2 Type II, FIPS 140-2, FIPS 197, and Common Criteria certifications. The platform supports SaaS, on-premises, private cloud, and hybrid deployment, with data residency options in North America and Europe. Administration is unified through the Enterprise Console, while end-user application access can be supported through a Chrome extension. Support is available in multiple languages through Broadcom's global Catalyst Partner network. This partner-led model provides broad regional reach, but support quality may depend on partner maturity and execution.

The roadmap includes native RDP session recording through ZTNA, connectorless ZTNA deployment through Google NCC, broader MCP support with enhanced AI-driven visibility, a Smart Assistant for policy management, expanded Shadow AI detection including semantically similar content, and browser-native DLP enforcement for Chrome, Edge, and Firefox.

Symantec SSE is best suited to mid-sized and large enterprises seeking a comprehensive, deeply integrated SSE platform with strong DLP and threat prevention across multiple access vectors. The combination of inline DLP, embedded RBI, and ZTNA in a unified data path is a strong differentiator for organizations with strict data governance and threat prevention requirements. Financial services, government, and regulated industries that need consistent policy enforcement across managed, unmanaged, and agentless scenarios will find the platform particularly relevant.


Strengths

- Uses Google's private backbone with direct peering to major cloud providers
- Inline DLP embedded natively in the ZTNA data path is among the deepest integrations in the market
- Embedded RBI eliminates the need for a separate web application firewall
- ZTNA as an MCP gateway with Secure On-Behalf-Of token exchange addresses the agentic AI identity-stripping problem
- Symantec DCS provides application control, host-based segmentation, and workload isolation for physical, virtual, and containerized data centers
- Broad global partner network across all major regions via the Catalyst Partner model, with localized time-zone and language support backed by Broadcom engineering

Challenges

- Data masking and redaction are not yet natively supported within DLP; support is on the near-term roadmap
- Workload identity federation is a long-term roadmap item
- Direct Broadcom support engagement is intentionally limited; customers are primarily served through regional distributors
- DLP capabilities within browser session controls are limited in comparison to other vendors

Cisco – Cisco Secure Access



Founded in 1984 and headquartered in San Jose, California, Cisco is a publicly listed global technology company with a portfolio spanning networking, security, collaboration, and observability. Its security business has expanded through multiple acquisitions, most recently Splunk in 2024. Cisco's Zero Trust offering is not a single product but a layered architecture combining Cisco Secure Access for SSE and ZTNA, Cisco Duo for identity and adaptive authentication, Cisco Identity Services Engine for network access control, Cisco Secure Workload and Hypershield for microsegmentation and workload protection, and Splunk for security analytics and operations. The primary product evaluated here is Cisco Secure Access, which combines ZTNA, SWG, CASB, Firewall-as-a-Service (FWaaS), VPN-as-a-Service, DLP, RBI, and AI Access controls in a cloud-delivered platform.

Cisco Secure Access is delivered as a cloud-native SSE platform with an identity-first design, where authentication and policy evaluation precede any connection to protected resources. ZTNA is implemented through an identity-aware relay architecture that hides private applications from the internet. Resource Connectors deployed as lightweight containers in customer environments, on-premises and in the cloud, establish outbound-only tunnels to the Secure Access cloud, eliminating inbound exposure. Agent-based access through Cisco Secure Client and clientless browser-based access are both supported. Clientless access covers web applications, SSH, and RDP.

A hybrid private access model allows local ZTNA enforcement to run directly on existing Cisco firewalls at branch or campus locations, keeping internal user traffic local rather than forwarding it to the cloud. SD-WAN integration with Cisco Catalyst, Meraki, and FTD automates tunnel provisioning and applies unified SSE policy for remote and office-based users. Traffic is encrypted using MASQUE/QUIC or TLS 1.3 with IPsec tunnel support for SD-WAN integration.

Access policies are managed through Cisco Cloud Control, a unified console covering internet, SaaS, and private application access. Policy is intent-based, with an AI policy assistant that translates natural language rules into enforcement configurations. Authorization incorporates identity from SAML providers such as Azure AD, Okta, and Ping, device posture from Cisco Secure Client or third-party MDM integrations, and contextual attributes like location, device type, and behavior. Integration with Cisco Identity Intelligence brings dynamic user trust levels into access policy to block or require re-authentication when behavioral or posture signals change. JIT and time-bound access are supported through ITSM integration. Approval workflows, policy verification, and a troubleshooting assistant are available for administrators.

Context and posture evaluation relies on Secure Client posture data, MDM compliance signals, identity risk from Cisco Identity Intelligence, and behavioral analytics from the CASB layer. Posture requirements can be differentiated by application type and user group. Missing posture signals reduce trust scores rather than causing instant denial, allowing gradual restrictions to be applied.

Segmentation in the broader Cisco Zero Trust architecture extends beyond the SSE layer. Cisco Secure Workload provides microsegmentation for data center and cloud workloads using agent-based and agentless enforcement. It discovers workload communication patterns and generates segmentation policies enforced through host firewalls, data processing units (DPUs), network firewalls, and cloud security controls. Cisco Hypershield adds kernel-level enforcement with eBPF for real-time workload segmentation. Both integrate with Cisco ISE for identity-enriched policies and with Cisco Cloud Control for unified management. Within Secure Access, lateral movement prevention is enforced through application cloaking and tunnels that prevent users from seeing or reaching resources beyond their entitlements.

Monitoring and analytics are centralized in a unified dashboard providing real-time visibility into sessions, flows, policy decisions, and identity context. The platform collects DLP events, threat inspection outcomes, UEBA behavioral signals, and standard telemetry. Administrators can trace policy decisions back to contributing signals in the console. Log export to enterprise SIEMs is supported through APIs and direct Splunk integrations, enabling correlation with endpoint and network telemetry across the Cisco portfolio. Automated responses include session termination, step-up authentication, account blocking,

and device isolation. Talos threat intelligence and AI-driven analysis enable both real-time detection and forensic investigations.

TLS inspection, protocol-level enforcement through FWaaS deep packet inspection, and RBI for risky or uncategorized sites are available inline within the Secure Access data path. Talos intelligence is also exposed through the Investigate API for automated threat investigation. AI Access controls provide visibility and enforcement for generative AI applications, including prompt injection detection, source code controls, and supply chain risk management for model repositories. MCP policy enforcement and intent-aware inspection of agent interactions address the identity-stripping problem in agentic AI systems.

DLP in Cisco Secure Access covers inline inspection of web, private application, SaaS API, and endpoint channels under unified policy management. Cloud malware detection scans files at rest in popular cloud services. Browser session controls through RBI, integrations for Microsoft Edge for Business, Chrome Enterprise Browser, and Island include copy-paste prevention, watermarking, file transfer controls, and keyboard controls for unmanaged devices. Data controls can be applied selectively based on user, device, and contextual signals.

Cisco Secure Access applies Zero Trust controls to NHIs through ZTNA and AI Access capabilities, governing API-based machine access, service accounts, and AI agents through the same policy layer used for human users. Machine and agent actions are logged and auditable. Agentic AI access through MCP is controlled by ZTNA acting as an MCP gateway. Cisco Identity Intelligence provides agent discovery and risk profiling, while workload identity and segmentation controls are handled through Cisco Secure Workload.

Cisco holds a broad certification portfolio including ISO 27001, SOC 2 Type II, PCI DSS, Common Criteria, FIPS 140-2, FIPS 197, and Germany C5. Cisco Secure Access offers a FedRAMP-certified package for US government environments. The SSE platform is cloud-delivered, with hybrid enforcement options through on-premises Cisco firewalls. Administration is unified through Cisco Cloud Control for multi-product management. Cisco Cloud Control entered controlled availability in the United States in June 2026, with global availability to follow soon. Support is available in multiple languages with 24x7 global coverage, while documentation is available in English. Talos Incident Response is available as an optional add-on.

The roadmap includes expanded MCP visibility and policy enforcement, native RDP session recording through ZTNA, intent-based policy management for Hybrid Mesh Firewall, further Splunk AI capabilities for agentic SOC use cases, and continued expansion of Cisco Identity Intelligence across SSE and access layers.

Cisco Secure Access is best suited to mid-sized and large enterprises seeking a comprehensive, integrated security architecture spanning SSE, identity, network segmentation, and endpoint protection. The depth of integration across Cisco's portfolio and the breadth of compliance coverage make it particularly relevant for regulated industries such as financial services, healthcare, and government.

Strengths

- Hybrid private access enforcement model allows ZTNA policies to run locally on existing Cisco firewalls


- Cisco Identity Intelligence provides unified agent discovery, risk profiling, and user trust levels across SSE, endpoint, and network
- Multi-layer Zero Trust enforcement spanning SSE, identity, endpoint, network segmentation, and SIEM from a single vendor
- MCP policy enforcement with intent-aware inspection of agent interactions addresses AI agent identity stripping
- Extensive certification coverage including FedRAMP, Common Criteria, and FIPS validations

Challenges

- Full Zero Trust architecture requires multiple separately licensed products
- Lacks native session recording through ZTNA, which is on the near-term roadmap
- It can be difficult for buyers to identify the minimum viable set of products needed to implement a coherent Zero Trust architecture
- On-prem and hybrid enforcement still depends on hardware refreshes: customers with older network products may face additional expenses

Fortra – Fortra Cloud Security Platform

FORTRA™



Leadership



Established in 1982 and rebranded as Fortra in 2022, the company is headquartered in Eden Prairie, Minnesota, and provides a broad portfolio spanning technology and cybersecurity. Its cybersecurity business is organized across offensive and defensive security capabilities, the latter including Zero Trust, data classification, DLP, digital rights management, vulnerability management, integrity and compliance monitoring, secure file transfer, and security awareness training. The evaluated solution is Fortra Cloud Security Platform, an SSE-based offering combining ZTNA, SWG, CASB, Firewall as a Service, DLP, secure email gateway, and DSPM in a cloud platform.

Fortra Cloud Security Platform is delivered as a cloud-native SaaS service built on AWS infrastructure and organized into three logical planes. The core service handles management, configuration, and log collection, with regional deployments supporting data

residency requirements. A globally distributed edge service layer acts as the data plane and enforcement point, while a global exchange layer synchronizes policy from the core service to edge nodes. Users connect to the nearest edge node through endpoint agents for Windows, macOS, and iOS, IPsec or SD-WAN tunnels for branch connectivity, and a reverse proxy model for unmanaged and BYOD access. Traffic is TLS 1.3 encrypted in transit, with IPsec available for site-to-site connectivity. The platform enforces authenticate-before-connect for traffic flows, with private applications exposed only through outbound connectors. Built-in application discovery enables VPN-to-ZTNA migration by giving administrators visibility into private applications across on-premises and cloud environments before policies are defined.

Access policies are organized across three enforcement layers that apply consistently to websites, SaaS, and private enterprise applications. At the network layer, FWaaS policies govern TCP and UDP ports and protocols using context such as identity, device posture, source IP, location, or user risk. At the TLS layer, domain and SNI inspection apply URL category and SaaS controls without requiring full decryption. At the application layer, full TLS inspection enables DLP, malware, and behavioral anomaly detection on payload content. A single unified policy model applies across traffic types. Application intelligence extends to the semantics of major SaaS platforms, including Microsoft OneDrive folder and file-level granularity, SharePoint site-level enforcement, Teams channel distinctions, and Slack workspace controls. This allows organizations to permit corporate instances while blocking personal or shadow ones. Contextual re-evaluation mid-session enables step-up MFA challenges, user coaching, session termination, and adaptive access restrictions as risk signals change.

Context and posture evaluation draws on identity attributes from SAML-compatible IdPs, device posture from UEM and MDM platforms, endpoint checks from EDR platforms, and IP reputation feeds from threat intelligence partners. User risk scores are computed using ML-based behavioral analysis across the data lake, incorporating DLP violations, download volumes, geo-anomalies, and activity baselines. Device and IP risk scores from external partners are combined with native signals. Posture requirements can be differentiated for each resource type, while missing posture signals result in fail-safe restricted access.

ZTNA enforces least-privilege access to private applications through outbound connector isolation, with each application explicitly published and scoped to authorized user groups and matching device posture. Network-level segmentation within private access restricts users to specific application instances, while application cloaking prevents lateral discovery. Inline malware scanning helps prevent malicious content from reaching endpoints or propagating to other workloads. Workload-to-workload and east-west segmentation inside data centers or cloud environments are not native SSE capabilities and require complementary tooling from Fortra's broader portfolio or third-party solutions.

Monitoring and analytics are unified across all traffic channels in a single data lake. The investigation interface allows administrators to query a user's complete activity history, including DLP violations, anomalous behaviors, application access patterns, and policy events. Shadow IT discovery, SaaS application risk profiling, and private application traffic visibility are available from the same console. Behavioral baselining and anomaly detection

run natively using ML models. Standards-based log export is supported to SIEM and XDR platforms. Bidirectional REST API integration allows external platforms to send risk signals into the policy engine and receive user risk score updates for coordinated enforcement.

Threat prevention includes inline URL and content filtering, web categorization, and a SaaS application intelligence engine that profiles applications against compliance posture and breach history. Malware scanning, TLS inspection, and Layer 3/4 firewall controls are available inline, including file traffic on private enterprise application connections. RBI is available through an OEM integration with Authentic8.

Data-centric security is the platform's clearest differentiator. The combination of inline DLP, Digital Rights Management (DRM), and data classification gives Fortra enforcement capabilities beyond many SSE platforms. Content actions include automatic file encryption, redaction, automatic application of classification labels from Microsoft Purview, Google, or Fortra Data Classification, user-specific watermarking, masking, quarantine, and document expiration. DRM enforcement allows files downloaded from private applications to remain protected after leaving the platform, with access controls, encryption, and expiration policies traveling with the file regardless of where it is later stored or shared. CASB API mode extends data protection to data at rest in cloud repositories. Browser-level controls are not natively provided within the SSE platform itself.

NHI protection is limited. Machine-initiated access to protected resources can be governed through the same ZTNA and policy enforcement framework used for human users, with actions logged and auditable. DLP enforcement applies to GenAI application traffic, including controls over prompt content and file uploads to sanctioned and unsanctioned LLMs. MCP server communication controls and agentic AI enforcement are on the near-term roadmap.

The platform holds ISO 27001, SOC 2 Type II, Common Criteria, and FIPS 197 certifications. Deployment is cloud-delivered, with no on-premises or self-managed option for the SSE platform. Administration is performed using a single console with RBAC-controlled delegation, and the end-user experience is transparent for agent-based deployments. Other products from Fortra's portfolio interoperate with the Cloud Security Platform through shared telemetry and coordinated policy. Support is available only in English.

The roadmap focuses on tenant-level controls for Microsoft Copilot, agentic AI and MCP server enforcement, endpoint DLP consolidation into the unified platform, and expanded AI-based content classifiers.

Fortra Cloud Security Platform is best suited to organizations that place data protection at the center of their Zero Trust strategy. The native integration of digital rights management into the ZTNA data path, combined with granular application intelligence and a broad range of content actions, makes the platform relevant for organizations with complex data sovereignty, compliance, or insider risk requirements.


Strengths

- DRM integrated directly into the ZTNA data path allows downloaded files to remain protected after leaving the platform
- Native DSPM capability provides visibility and governance over sensitive data at rest in cloud repositories
- Automatic application of classification labels from third-party engines with corresponding DRM enforcement on download
- Deep application intelligence covers SaaS platforms at granular levels, enabling precise corporate and personal instance separation
- Bidirectional risk signal exchange through REST API with external platforms
- Broad data protection capabilities, including automatic encryption, redaction, watermarking, masking, quarantine, and expiration

Challenges

- Cloud-only delivery limits applicability for air-gapped or sovereignty-constrained environments
- NHI protection is limited; MCP server communication controls and agentic AI enforcement are on the roadmap
- Browser-level controls rely on OEM RBI integration
- East-west workload-to-workload segmentation requires third-party tooling
- English-only support is a constraint for global enterprise deployments

GoodAccess – GoodAccess ZTNA Platform



Founded in 2014 and headquartered in Prague, Czech Republic, GoodAccess is a privately held company owned and operated within the European Union (EU). Their EU focus may be a strategic differentiator for organizations with data residency and regulatory requirements, including those driven by NIS2. It offers a single product, the GoodAccess ZTNA Platform, delivered as a SaaS service. The platform is designed primarily for small and midsize enterprises.

GoodAccess ZTNA is built on a centralized control plane with distributed enforcement through a global gateway layer. Each customer receives dedicated gateway infrastructure and dedicated IP addresses, distinguishing the platform from shared architectures commonly used by larger SSE vendors. Connectivity between users and protected resources is carried through encrypted tunnels, with WireGuard as the default protocol and IPsec and OpenVPN also supported. Traffic and management communications are protected using industry-standard encryption, including TLS 1.3 for control plane, management, and API communications. The platform enforces authenticate-before-connect for protected resources, and protected systems have no exposed inbound ports.

Private network connectivity is handled through Cloud & Branch Connectors, which establish standards-based IPsec site-to-site tunnels between GoodAccess gateways and whatever equipment the customer already has: existing edge routers, cloud VPN gateways, or virtual machine instances. The platform does not require its own physical or virtual appliance inside the customer environment. For mid-market customers, this means a faster path off legacy VPNs without taking on additional vendor-specific hardware or software footprint. User access requires an endpoint agent; clientless or browser-based access is not currently supported. Agentless access for unmanaged devices is on the roadmap.

Access policy is based on identity-driven microsegmentation using a structured abstraction called virtual access cards. Instead of requiring administrators to write firewall rules, access cards define logical relationships between users or groups and named resource groups defined by IP address or FQDN, port, and protocol. Enforcement rules are generated from these definitions, reducing the likelihood of policy conflicts and sprawl. This model is one of the platform's important design decisions: it trades some advanced policy granularity for operational clarity, useful for smaller IT teams.

Access cards support time-based restrictions, including working hours, day-of-week rules, and expiration dates, making them useful for contractor and temporary access. Geographic restrictions can also limit access based on the user's detected location. Policy granularity reaches the server, application, and service level, but not the API route or function level. Risk-based access and approval workflows for sensitive access requests are not currently available and are planned for a later roadmap phase. Step-up authentication can be enforced by terminating the session via API and requiring the user to reconnect with MFA or biometrics.

Context and posture evaluation is performed by the endpoint agent at session establishment and periodically during active sessions. Posture checks include OS version compliance, Endpoint Protection and EDR presence, disk encryption, lock screen enforcement, and custom checks for registry keys, files, services, or processes. Device-specific policies can be defined separately for Windows, macOS, Linux, Android, and iOS. If a device fails a posture check, access is denied, and the user sees a configurable message. Native integrations with MDM and EDR platforms are not yet available. Deeper integrations and more granular per-resource posture requirements are on the roadmap.

Segmentation is identity-driven and independent of network topology. Policies are defined against named microsegments rather than VLANs or broad IP ranges. The default-deny model at the gateway layer blocks communication unless it is explicitly permitted by an active access card. East-west traffic control is partially supported when at least one side of the communication runs the GoodAccess agent. Native workload-to-workload segmentation is not yet supported. Local gateways for on-premises environments, intended to extend east-west controls to internal network segments and OT environments, are planned.

Monitoring capabilities cover session-level logs per gateway, access logs showing user sessions against microsegments with duration and data volumes, device posture histories, and administrative audit logs for configuration changes. Active session visibility is available in the admin console. Logs can be exported through syslog or REST APIs, allowing

integration with SIEM platforms such as Microsoft Sentinel, Splunk, and IBM QRadar. GoodAccess delegates behavioral detection and correlation to external SIEM and SOAR tools. These tools can trigger enforcement actions through API, including session termination, user blocking, and device isolation. The console also provides geographic access maps showing where users connected from.

The combination of authenticate-before-connect, default-deny enforcement, device posture, and application cloaking reduces exposed attack surface and limits unauthorized discovery. DNS-layer filtering uses external threat intelligence to block known malicious and phishing domains, and customers can define additional block lists. However, GoodAccess does not currently provide inline TLS inspection or malware scanning. A Secure Web Gateway capability with local agent-based proxy and packet inspection is in development and is expected to extend threat prevention into web traffic content.

Data-centric security is not currently implemented. Data access is governed at the network and session level through microsegmentation policies. Some of these capabilities are expected to emerge with the planned SWG component, but organizations with immediate requirements for data loss prevention or content-aware controls will need complementary solutions.

Machine identities connected through the agent-based model are subject to the same access and posture controls as users, and their actions are logged and auditable. However, advanced features such as service account lifecycle management or controls for AI agents are not currently supported. Expanded NHI and workload segmentation capabilities are part of the longer-term roadmap toward a broader SSE platform.

GoodAccess holds ISO 27001 certification and SOC 2 Type II attestation covering security, availability, and confidentiality. Deployment is SaaS, with no standard on-premises or self-managed option. Tenant data and logs are stored within the EU. The administrative experience is a clear differentiator. The virtual access card model allows IT generalists without firewall expertise to deploy and manage Zero Trust access, while the visual policy representation makes access relationships easy to understand and verify. The REST API enables programmatic management and integrations with external security tools. Support is provided in English and German, with AI-assisted translation available for other languages. Documentation is available in English.

The roadmap focuses on SWG with local proxy-based inspection, local gateways for east-west and OT network coverage, dynamic context-based access control, expanded posture signals, and EDR and identity risk integration. The longer-term vision is a unified identity, network, and workload security platform with broader non-human identity and SSE capabilities.

GoodAccess ZTNA is well suited to small and midsize organizations seeking a straightforward Zero Trust network access solution that does not require specialized security engineering. Its dedicated gateway model, EU data residency, standards-based connectivity, and clean administrative experience make it especially relevant for European organizations with sovereignty requirements and limited IT resources.


Strengths

- Each customer receives dedicated gateway infrastructure and dedicated IP addresses
- EU-headquartered and EU-hosted, addressing NIS2, GDPR, and data residency requirements for European organizations without complex configuration
- Virtual access card model translates complex firewall rule authoring into a visual, structured abstraction without the need for specialized security expertise
- ISO 27001 and SOC 2 Type II certifications are notable for a company of its size and market segment
- WireGuard as the default protocol provides modern, high-performance encryption with lower overhead than legacy VPN protocols
- Support available in German in addition to English, relevant for the primary DACH market

Challenges

- No agentless or browser-based access is currently available. BYOD and unmanaged devices are supported through the endpoint agent, which is required for user connections
- Web threat prevention is currently limited to DNS-layer filtering; inline SWG inspection capabilities are in development
- Native EDR and MDM platform integrations are not yet available
- Risk-based access, approval workflows for sensitive access requests, and east-west workload segmentation are not currently supported
- Go-to-market focus and primary market presence remain Europe-centric; the customer base is primarily SMB and mid-market

Illumio – Illumio Breach Containment Platform



Founded in 2013 and headquartered in Sunnyvale, California, Illumio focuses exclusively on Zero Trust segmentation and breach containment. Its platform is designed to complement, not replace, ZTNA and SSE solutions: where those primarily govern north-south user access, Illumio concentrates on east-west lateral movement control, workload microsegmentation, and rapid containment after compromise. The evaluated offering is the Illumio Breach Containment Platform, consisting of Illumio Segmentation for visibility, policy management, and enforcement across hybrid and multi-cloud environments, and Illumio Insights, a cloud detection and response capability powered by AI.

The platform uses a two-tier architecture. The Policy Compute Engine (PCE) acts as the control plane, while enforcement is performed by Virtual Enforcement Nodes (VENs), cloud-native controls, container controls, or network enforcement points. The PCE is available as a SaaS service with regional options in the US, EMEA, and APAC, or as a self-managed customer deployment. A defining characteristic is that Illumio does not directly intercept traffic. Instead, the PCE sends policy-based instructions to native enforcement points, such as iptables on Linux or Windows firewalls. This avoids latency overhead and throughput bottlenecks while allowing enforcement to continue even when the control plane is temporarily unavailable.

For cloud workloads, enforcement is applied through cloud-native controls such as AWS Security Groups, Azure NSGs, GCP firewall rules, and cloud-native firewalls such as Azure Firewall. Container enforcement is supported through Kubernetes Network Policy, enforced via CNIs such as Cilium and Calico. Agentless enforcement for data center workloads and OT devices is available through network enforcement points such as load balancers and switches. Agent coverage spans Windows, macOS, major Linux distributions, Solaris, and AIX, including legacy operating systems that remain common in large enterprises. Specialized environments are addressed through OEM-style partnerships: IBM Z mainframes and IT/OT industrial networks are supported via partner integrations rather than native agents.

Illumio instantiates policies through component labeling. Labels describe workloads in business terms such as application, environment, role, business unit, and location. Policies thus remain valid as IP addresses, cloud infrastructure, and network topology change. The Illumio Console provides a traffic map showing observed communication flows between labeled workload groups, a policy generator that recommends rules from observed dependencies, policy templates for common use cases such as ransomware containment and environment separation, and draft mode for visualizing policy impact before enforcement. Policy versioning allows administrators to stage and roll back changes safely.

Organizations can begin to assess and plan in visibility-only mode and then proceed to enforcement mode. Policy management is also available through REST APIs, a Terraform provider, Python SDKs, and policy-as-code integrations for DevOps pipelines. RBAC can be delegated to application owners scoped to their own workload groups, supporting decentralized policy management at scale.

Context and posture evaluation is extensive for workload environments. The platform computes vulnerability exposure based not only on whether a vulnerability exists, but whether the affected port is actually reachable from other workloads. EDR, XDR, and SIEM integrations (including CrowdStrike Falcon, Splunk, Microsoft Sentinel, and Palo Alto Networks Cortex XSOAR) can drive automated quarantine decisions based on detections. Identity-aware policies use Active Directory, Entra ID, and LDAP group membership to enforce workload access controls tied to user identity, including the ability to distinguish different users on the same endpoint. Domain awareness allows policies to become more restrictive when devices move away from corporate networks.

Segmentation is Illumio's core capability and strongest area of differentiation. The label-based model applies consistently across physical servers, virtual machines, containers, cloud workloads, endpoints, OT, and IoT environments. Application dependency maps built from observed traffic help organizations understand communications before writing policy. Workload-level secure connectivity is supported through Secure Connect, which uses IPsec-based encryption for workload-to-workload channels, and through machine authentication using mTLS or Kerberos. Process-based segmentation can specify not only which workloads may communicate, but which processes on those workloads may initiate or accept connections. For Kubernetes environments, service mesh mTLS enforcement integrates with Kubernetes-native controls.

Monitoring and analytics are powered by the AI Security Graph behind both Illumio Segmentation and Illumio Insights. The platform collects flow telemetry across protected environments and enriches it with workload labels, threat intelligence, vulnerability data, and identity context. Insights Hub provides dashboards for traffic flows, firewall policy coverage, malicious IP communications, external data transfer monitoring, and shadow AI and LLM access detection. Behavioral baselining uses ML to establish normal traffic patterns per workload label and flag deviations. The Insights Agent acts as an AI-driven investigation assistant, analyzing telemetry, identifying significant threats and anomalous patterns, mapping findings to MITRE ATT&CK techniques, and recommending remediation actions. Findings can be sent to ServiceNow, and telemetry export is supported through syslog and integrations with Splunk, Microsoft Sentinel, and Google SecOps.

In full enforcement mode, default-deny segmentation can block lateral movement even when credentials are valid, because compromised accounts cannot reach workloads outside explicit policy. Real-time flow visibility and one-click quarantine support rapid containment of ransomware and other spreading threats. Illumio Insights identifies command-and-control communications, exfiltration patterns, and risky service traffic through behavioral analytics and threat intelligence enrichment. Illumio does not inspect content or enforce DLP policies. However, it contributes to data governance through visibility into cross-region and inter-cloud flows, external data transfer monitoring, shadow AI and LLM access detection, and data residency enforcement through segmentation policy.

Workloads, APIs, service accounts, and AI systems can be treated as first-class policy subjects using the same label-based model. Machine identities are authenticated through mTLS, Kerberos, or tokens with configurable expiry and rotation. Workload identity federation is supported, and least-privilege enforcement for automated systems and AI pipelines uses the same policy engine as other workloads.

Illumio holds FedRAMP, Common Criteria, SOC 2 Type II, and ISO 27001 certifications. Deployment options include SaaS, self-managed, and hybrid models. The end-user experience is completely transparent, since VENS operate without user interaction. Administration is mature, with a comprehensive console, APIs, Terraform support, SDKs, and policy-as-code capabilities. Support and documentation are available in multiple languages.

The roadmap includes identity risk and blast radius visualization through Entra ID correlation, expanded anomaly detection, additional AI assistants for security operations and compliance, agentless Zero Trust enforcement through Check Point firewalls and NVIDIA DPUs, deeper CrowdStrike integration for endpoint segmentation, and broader Kubernetes CNI support.

Illumio Breach Containment Platform is best suited to mid-market and large enterprises that need to contain lateral movement and enforce workload-level segmentation across complex hybrid and multi-cloud environments. Its strengths lie in workload coverage, policy scale, host-native enforcement, application dependency visibility, and operationally safe rollout. It is particularly relevant for organizations with large data centers, legacy infrastructure, critical workloads, or complex Kubernetes environments.


Strengths

- Label-based policy model defines segmentation in business terms, making policies topology-independent and portable
- Process-level segmentation, with policies enforced on workload processes and ports rather than IP addresses alone
- Host-native enforcement via native OS firewall controls and cloud-native controls
- Proven in massive-scale customer deployments with millions of workloads
- Illumio Insights Agent provides AI-driven investigations, recommends remediation
- Shadow AI and LLM access detection, cross-region and inter-cloud data transfer monitoring, and data residency enforcement
- Strong compliance posture with certifications in all major geographic areas and industries

Challenges

- Must be paired with separate remote access and internet enforcement tools for complete Zero Trust coverage
- Continuous re-evaluation of active sessions is not supported
- Inline traffic inspection is absent by design
- Policy-as-code is not natively supported out of the box; customers must build their own GitOps integrations
- Pre-built, tested connectors for third-party security tools remain limited compared to leading peers, though the ecosystem is actively being expanded

Ivanti – Ivanti Neurons for Zero Trust Access



Leadership



Founded in 2017 through the merger of LANDESK and HEAT Software, and subsequently expanded through several acquisitions, Ivanti is a privately held company headquartered in South Jordan, Utah. The company operates globally across a broad IT and security portfolio that includes autonomous endpoint management, exposure management, IT service management, and network security. The product evaluated here, Ivanti Neurons for Zero Trust Access, is the company's dedicated ZTNA offering and a core component of its broader portfolio.

Ivanti Neurons for Zero Trust Access is built on a Software-Defined Perimeter architecture aligned with NIST SP 800-207. The platform consists of a cloud-hosted Controller acting as

the Policy Decision Point (PDP), customer-deployed gateways serving as PEPs, and the Ivanti Secure Access Client installed on user endpoints. The Controller runs exclusively on Microsoft Azure, with regional availability to support data residency requirements. User application traffic flows directly from the client to the customer-deployed gateway and does not traverse Ivanti's cloud infrastructure. The Controller receives policy metadata and authentication signals only, allowing customers to retain full sovereignty over application traffic.

Gateways are deployed as virtual appliances in customer environments close to protected resources and support VMware ESX, AWS, Azure, GCP, Oracle Cloud, KVM, and OpenStack. mTLS is used for connections among the client, Controller, and gateways. Client support covers Windows, macOS, Linux, iOS, and Android through a single unified agent that also supports Ivanti Connect Secure VPN under the same Unified Access License, with VPN and ZTNA connections running concurrently and ZTNA taking priority when available. Native agentless or browser-based access is not yet available, though HTML5-based access can be delivered as an add-on through Ivanti Connect Secure under the same license. Native support is planned as a longer-term roadmap item.

Access policy in Ivanti nZTA is based on four main parameters: user or group identity from identity providers such as Azure AD or Okta; device posture evaluated by the integrated Host Checker component powered by OPSWAT; gateway assignment; and application definition. Policies determine which users, on which compliant devices, may access which applications through which gateways. Authorization granularity extends from server and application level down to IP addresses, FQDNs, port ranges, and protocols including TCP, UDP, RDP, SSH, and non-standard ports. API route-level and function-level enforcement are not supported.

Time-bound access policies and configurable session duration limits are supported. Continuous session re-evaluation is implemented through Continuous Adaptive Risk and Trust Assessment (CARTA), with device posture checked during active sessions and user risk scores recalculated accordingly. If a risk score crosses a configured threshold, the platform can terminate the session, require step-up authentication on the next login, issue a warning, or deny subsequent access. ABAC is supported through behavioral risk signals and contextual attributes such as location, network, and device compliance. JIT access can be implemented through time-based policies, although approval workflows for sensitive access requests are not currently available. A notable recent addition is the AI Gateway Migration Assistant, which analyzes existing Ivanti Connect Secure VPN configurations, translates legacy access rules into Secure Access Policies, and generates step-by-step migration plans.

Host Checker provides broad endpoint compliance assessment before and during sessions, using predefined and custom checks for endpoint security tools, OS and patch levels, encryption, files, registry keys, processes, ports, location, network state, and CVE exposure. The OPSWAT Endpoint Security Assessment Plug-in (ESAP)-based library supports many third-party endpoint security products, while integration with Ivanti Risk-Based Vulnerability Management adds context on vulnerable applications and OS versions based on known CVEs and exploit activity. Access restrictions can be applied automatically to devices

running vulnerable or non-compliant software, with posture requirements differentiated per application or gateway group through risk scores and administrator-defined thresholds. The limitation is that this is primarily a posture and vulnerability-risk model, not full third-party EDR or identity-risk ingestion. Ivanti can validate whether supported endpoint security tools are present and active, and ZTA integrates with Ivanti's own MDM and endpoint-management portfolio, but direct consumption of live external EDR signals, such as third-party device risk scores, appears limited. Similarly, external identity-risk signals such as risky sign-ins, compromised credentials, or impossible travel are not yet consumed directly for continuous access decisions.

Each secure access policy creates independent per-application micro-tunnels between the enrolled client and the designated gateway. However, east-west workload traffic, service-to-service segmentation, and workload-level microsegmentation are not supported today. Gateway groups provide high availability and load balancing, and multiple gateways can support different application groups within the same deployment without additional per-gateway licensing. Application Discovery, available for existing Ivanti VPN customers, surfaces observed application traffic from VPN sessions and can automatically recommend ZTNA policies, further supporting migration planning.

Monitoring and analytics are delivered through a centralized administrative dashboard with chart visualizations of user, device, gateway, and application flows. UEBA-based user risk scoring is a core analytical capability, combining device posture, vulnerability data, location, behavioral anomalies, and policy violations into a per-user risk score that can drive enforcement decisions. Administrators can trace risk scores back to contributing signals for investigation. Historical data retention is 30 days by default, with extended retention planned. Log export is supported through syslog and REST APIs for SIEM integration. Behavioral baselining and anomaly detection are available through the UEBA engine, and automated response options include session termination, step-up authentication on the next login, account blocking for subsequent logins, or warning messages.

Threat prevention is primarily delivered through the platform architecture: SDP controls, mutual TLS, posture enforcement, and UEBA-driven risk scoring. A Web Application Firewall (WAF) capability is on the near-term roadmap as part of the gateway migration to NGINX and is intended to protect applications hosted behind the gateway. NHI support is limited. Device certificates are issued to enrolled endpoints as part of the mTLS model and rotated automatically before expiry. Broader support for workloads, IoT devices, and AI agents is identified on the 24 to 48-month roadmap.

Ivanti holds ISO 27001 certification and SOC 2 Type II attestation across all five trust service criteria. FIPS 140-2 certification is not currently held, although migration of the gateway to Oracle Linux with FIPS 140-3 compliance is planned. Deployment is based on a SaaS-hosted Controller with customer-deployed gateways, supporting on-premises, private cloud, and public cloud placement. The Unified Access License includes both ZTNA and VPN gateway entitlements, simplifying the transition for organizations with existing Ivanti VPN infrastructure. Support is provided primarily in English, with best-effort coverage in German, French, Spanish, Chinese, and Japanese. Documentation is available in English only.

The roadmap includes clientless access, Digital Experience Monitoring, an AI-based threat neutralizer for runtime risk evaluation, post-quantum cryptography, inline traffic inspection, and unified visibility across Ivanti's other products.

Ivanti Neurons for Zero Trust Access is best suited to mid-market and SMB organizations transitioning from Ivanti Connect Secure VPN to ZTNA and looking for a managed, gradual migration path under a unified license and client. Its SDP architecture, strong endpoint posture assessment, UEBA-based risk scoring, and VPN coexistence capabilities make it a practical option for organizations that need reliable user-to-application access control without the operational complexity of a full SSE platform.


Strengths

- Unified Access License bundles both ZTNA and VPN under a single license with a common client, enabling gradual VPN-to-ZTNA migrations
- AI Gateway Migration Assistant provides a differentiated migration tool for the large installed VPN base
- User risk scoring incorporates device posture, vulnerability data, location, behavioral anomalies, and policy violations into a composite score that drives enforcement decisions
- Unique CVE-to-access-policy linkage enables automatic access restrictions for devices running software with known active exploits
- Application Discovery analyzes observed VPN traffic patterns and automatically recommends ZTNA policies

Challenges

- Browser-based access requires an add-on module, and native clientless access is a longer-term roadmap item
- Third-party EDR and external identity-risk integrations remain limited; access decisions rely more on posture checks than on live risk telemetry
- East-west workload segmentation, service-to-service communication controls, and workload-level microsegmentation are not part of the solution
- Limited support for NHIs
- Inline traffic inspection, WAF capabilities, and post-quantum cryptography are all on the roadmap, not yet available

Microsoft – Microsoft Security



Founded in 1975 and headquartered in Redmond, Washington, Microsoft is a global technology company operating across cloud computing, productivity software, devices, AI platforms, and enterprise security. Microsoft Security spans identity and access management, endpoint protection, cloud security, information protection and governance, threat intelligence, SIEM and SOAR, and network security. The product evaluated here is Microsoft Security as an integrated Zero Trust platform, with Microsoft Entra for identity and access control, Microsoft Intune for device management, Microsoft Defender XDR for threat protection and endpoint security, Microsoft Sentinel for SIEM and SOAR, Microsoft Purview for data governance and protection, Global Secure Access for network enforcement, and Microsoft Defender for Cloud for cloud workload protection.

Microsoft's Zero Trust architecture is centered on Microsoft Entra Conditional Access as the PDP. It combines signals from identity, device posture, network context, application sensitivity, behavioral risk, and threat intelligence to make access decisions across heterogeneous environments. The architecture separates control and data planes across a globally distributed cloud-native infrastructure. Enforcement is multi-modal: agent-based through Defender for Endpoint and Intune; agentless and browser-based through Entra Application Proxy and Defender for Cloud Apps session proxy; network-based through Global Secure Access, Azure Firewall, NSGs, and Private Link; embedded through service mesh integrations such as Istio; and token-based through Continuous Access Evaluation. Traffic is TLS-encrypted in transit (TLS 1.2 minimum, TLS 1.3 preferred), with IPsec for site-to-site connectivity and HTTP/3 over QUIC for modern Microsoft cloud services. Protected resources can be accessed through outbound connectors, reducing direct inbound exposure.

Conditional Access uses RBAC, ABAC, contextual and risk-based access, and continuous session re-evaluation. It enables near real-time enforcement when critical events occur, including account changes, password resets, elevated identity risk, location changes, and behavioral anomalies. Policy granularity ranges from server and application level (enforced by Conditional Access) down to API route, function, and method level via OAuth scopes, app roles, and Azure API Management. JIT and time-bound privileged access is delivered through Entra Privileged Identity Management, with approval workflows, MFA requirements, business justification, and automatic expiry. Step-up authentication can be triggered when risk or sensitivity increases. Conditional Access What-If Analysis supports policy simulation, and policy insights provide optimization guidance.

Microsoft ingests device intelligence from Intune and Defender for Endpoint risk levels; identity risk from Entra ID Protection, including impossible travel, leaked credential detection, atypical behavior, and sign-in anomalies; cloud posture from Defender for Cloud; and network context such as IP, location, and trusted network changes. External signals from tools such as CrowdStrike, Jamf, and other MDM and EDR platforms can feed into Conditional Access through Intune device compliance and Mobile Threat Defense connectors, while their alerts and incidents flow into Microsoft Sentinel and Microsoft Defender via the Microsoft Graph Security API. Posture requirements can drive graduated restrictions such as session-only access, restricted downloads, or browser-only access, rather than simple allow or block decisions.

Segmentation spans identity, network, and workload layers. Conditional Access enforces least privilege independently of network topology. Azure Virtual Networks, NSGs, ASGs, Private Link, and Azure Firewall provide workload and application isolation. Global Secure Access covers web and SaaS traffic filtering and identity-aware connectivity to private applications and resources, supporting VPN replacement. Workload-to-workload segmentation is enforced through workload identities, mTLS, managed identities, and service mesh integration on AKS. East-west Kubernetes traffic is governed through Kubernetes Network Policies and Istio/Linkerd service mesh mTLS with workload identity. Microsoft's microsegmentation approach is therefore multi-layered rather than centered on a single dedicated product.

Monitoring and analytics are centralized through Microsoft Sentinel, which functions as the data lake, SIEM, and SOAR platform across the security stack. Sentinel ingests telemetry from Microsoft security products and third-party sources through prebuilt connectors and formats such as CEF and syslog. UEBA builds behavioral baselines for users, service accounts, and workloads. Defender XDR enables incident investigation across identity, endpoint, email, cloud, and network signals in a single incident graph, with AI-assisted triage. Microsoft Security Copilot helps analysts investigate incidents, query threat intelligence, interpret policies, and generate playbooks. Administrators can trace access decisions back to contributing signals through Conditional Access Insights.

Threat prevention is integrated across identity, endpoint, email, cloud, network, API, and AI layers. Identity protection covers account compromise detection, token theft detection, impossible travel, and automated high-risk user blocking. Defender for Endpoint provides malware scanning, attack surface reduction, behavioral detection, and network protection. Defender for Office 365 adds URL detonation, Safe Attachments, and Safe Links. Defender for Cloud Apps applies inline protocol inspection for HTTP/S session anomalies and risky actions. Azure Firewall Premium provides TLS inspection and protocol validation. For AI workloads, Azure AI Content Safety Prompt Shields detect and block jailbreak and prompt injection attempts against Azure AI Foundry-hosted models, complemented by Defender for Cloud AI threat protection. Microsoft threat intelligence is informed by broad telemetry across the Microsoft ecosystem.

Microsoft Purview provides data protection capabilities, including sensitivity label-based classification and protection, automated labeling, DLP across endpoints, email, SaaS applications, AI interactions, and insider risk management. Sensitivity labels can travel with content across Microsoft 365 and can be referenced by Conditional Access and session controls to block downloads, prevent printing, apply dynamic watermarks, or enforce browser-only access. DLP policies can apply inline through Defender for Cloud Apps session controls, with cloud app discovery and unsanctioned app blocking extended through integrations with secure web gateways such as Zscaler. Data residency and sovereignty controls allow customers to specify regions for tenant data and logs, with customer-managed key options available.

Agent 365 is Microsoft's registry and governance platform for AI agents, providing discovery, identity assignment through Entra Agent ID, behavioral observability, risk signal integration, lifecycle management, and orphaned agent detection. Agent ID enables Conditional Access policies for agents, including real-time risk assessment and blocking of high-risk agents. NHIs, including workloads, APIs, service accounts, containers, Kubernetes pods, and automation pipelines, are governed through Entra Workload ID, managed identities, and workload identity federation with cross-cloud support.

Microsoft maintains a broad certification portfolio including ISO 27001, SOC 2 Type I and Type II across all five trust service criteria, FedRAMP High, PCI DSS (for selected cloud services), Germany C5, UK Cyber Essentials, and FIPS 140-2. The platform is delivered primarily as a cloud service, with on-premises coverage through connectors, agents, and Azure Arc. Support is available in major languages, with extensive documentation, training, assessment tools, workshops, and a broad partner and integration ecosystem.

The near-term roadmap includes broader availability of Agent 365 registry and access control capabilities, Conditional Access for agents through Agent ID, runtime DLP for prompts in Copilot Studio and Azure AI Foundry agents, shadow MCP server discovery, unified risk scoring across Entra and Defender, predictive identity threat shielding, Sentinel data federation using Microsoft Fabric, and expanded AI posture management through Defender for Cloud.

Microsoft Security is best suited to organizations already invested in the Microsoft ecosystem and seeking to implement Zero Trust through a deeply integrated platform rather than assembling capabilities from many vendors. The platform is particularly relevant for regulated industries, government, large enterprises, and organizations adopting agentic AI at scale.


Strengths

- Agent 365 registry provides the most complete agentic identity governance framework evaluated
- Conditional Access as a unified PDP integrates signals from identity, device, network, application, behavioral risk, and threat intelligence across Microsoft and third-party sources
- Microsoft Purview delivers sensitivity label-based classification, inline DLP, insider risk management, and data governance
- Entra Workload ID and Entra Agent ID provide broad governance across non-human identities, including workloads and AI agents
- Prompt Shield detection for jailbreak attempts and shadow MCP server discovery address AI-specific threats at scale
- Broadest compliance certification portfolio across all evaluated vendors

Challenges

- The Zero Trust architecture requires many separately licensed and managed Microsoft products, which may create significant complexity and cost for customers
- Network-level ZTNA enforcement via Global Secure Access is newer to the market than established ZTNA specialists, with deepest adoption in Microsoft-stack environments
- Organizations not using Azure or Microsoft 365 as primary platforms will find limited applicability for much of the Zero Trust enforcement layer

NetFoundry – NetFoundry



Founded in 2018 and headquartered in Charlotte, North Carolina, NetFoundry is the commercial developer and maintainer of OpenZiti, the open source Zero Trust networking software that forms the technical foundation of its platform. NetFoundry is offered as a subscription platform with vendor-hosted SaaS, self-hosted on-premises or private cloud, and hybrid configurations available across management, control, and data plane components. The NetFoundry platform primarily addresses the use cases of workload connectivity, microsegmentation, API security, and machine identity, with human-centric remote access supported as a secondary capability via NetFoundry Frontdoor. OEM integrations and embedded SDKs allow NetFoundry capabilities to be embedded directly into cloud-native, Kubernetes, AI, and networking applications and infrastructure.

NetFoundry is built on a strict three-plane architecture. The management plane provides centralized administration, multi-tenancy, identity lifecycle management, policy definition, and API access. The control plane acts as the policy decision point, authenticating identities, evaluating access policies and posture requirements, issuing certificates, and continuously authorizing sessions. The data plane consists of distributed software-based edge routers that form a private, per-customer overlay fabric mesh. Smart routing selects low-latency paths across available underlay links, with automatic rerouting on link or node failure. Edge routers operate in a redundant mesh topology and support active-active high availability.

A defining architectural characteristic is that private keys are generated on the originating endpoint and never leave it. The control plane issues certificates but has no access to private key material, meaning NetFoundry cannot decrypt customer traffic. Connection options include lightweight agents for Windows, macOS, Linux, iOS, and Android; embedded SDKs for Golang, C, Java, .NET, Swift, and Python; Kubernetes sidecar and DaemonSet deployment; containerized edge routers; and NetFoundry Frontdoor for agentless browser access. The platform can also be deployed through OEM integrations and supports air-gapped and edge environments.

NetFoundry follows the ABAC model and evaluates identity attributes. Authorization granularity extends from server and application level down to API route. Policies are centrally managed through the console, REST APIs, or Terraform. NetFoundry performs continuous access evaluations: policy changes, certificate revocation, or posture check failures immediately terminate active sessions and make services invisible to the affected identity. Risk scoring itself is not performed natively by NetFoundry and must be sourced from a third-party EDR, XDR, or identity platform. Once integrated via API, those external risk signals can trigger policy changes, session termination, or re-authentication in real time.

Device posture evaluation, which is distinct from risk scoring, is performed natively by the NetFoundry endpoint agent and SDK. The agent collects deterministic device-state signals (OS version, domain membership, MAC address, running processes, MFA state, and certificate validity), and these checks are evaluated continuously during sessions, not just at connection time. In addition to these native checks, posture attributes from external MDM, UEM, and EDR sources can be mapped into policy.

Posture requirements are configurable per service, allowing different resources to enforce different context conditions. Workload runtime context, including identity, deployment method, certificate attributes, and SDK usage, are also available for evaluation in authorization decisions. Native behavioral baselining and anomaly detection are not provided; these functions require external SIEM and SOAR tools.

Segmentation is the area where NetFoundry diverges most clearly from conventional ZTNA products. The overlay model enforces service-level microsegmentation independent of network topology, infrastructure ownership, or organizational boundary. The same policy model governs both human user access and machine-to-machine traffic between workloads, services, and APIs, spanning cloud, on-premises, edge, OT, IoT, and partner environments. Because connectivity is established only between authenticated identities and named services, lateral movement is prevented by design.

Kubernetes-native enforcement is supported through pod-level sidecars, container-level SDKs, and cluster-level edge routers. OT environments are addressed through NetFoundry zLAN, which uses eBPF within the edge router to extend identity-based segmentation to devices that cannot run a NetFoundry endpoint (such as Programmable Logic Controllers [PLCs], controllers, and IoT sensors). zLAN operates at the IP layer and is protocol-agnostic; it does not perform deep packet inspection of OT protocols.

Monitoring and analytics cover session-level logs per identity, L3/L4 flow metadata, policy decision records, certificate lifecycle events, posture check results, and control plane and data plane health metrics. Real-time visibility into active sessions, identities, services, and connection status is available through the management console. Telemetry can be streamed through syslog or REST APIs. Automated response actions available through APIs or direct policy enforcement include session termination, identity revocation, step-up authentication enforcement, and device isolation.

The management API is itself protected behind a NetFoundry overlay network. NetFoundry does not perform TLS termination, malware scanning, or protocol inspection, but traffic can be routed through external security tools where required. Token theft and session hijacking are mitigated through several layered controls. Private keys are generated on the endpoint and never leave it (NetFoundry refers to these as sovereign private keys), so a stolen bearer token cannot be replayed from another device without the original key. Session keys are ephemeral and bound to mTLS, and authorization is re-evaluated continuously, so even a hijacked session is short-lived and revocable.

NetFoundry supports workloads, APIs, service accounts, automation pipelines, IoT and OT devices, and AI agents. NHIs receive X.509 certificates with automated lifecycle management through NetFoundry's built-in certificate authority. Workload identity federation is supported through external PKI, SPIFFE/SPIRE, and third-party CAs. NetFoundry MCP Gateway and LLM Gateway extend this model to agentic AI by making LLM endpoints and MCP tool servers dark by default unless an agent is explicitly authorized. Agents receive their own identities instead of inheriting API keys or service account credentials. The gateway layer also provides AI token tracking and LLM load balancing.

NetFoundry holds SOC 2 Type II attestation across all five trust service criteria. FIPS 140-2-compliant mTLS is supported for relevant government and defense use cases. Deployment options include vendor-hosted SaaS, self-hosted on-premises or private cloud, and hybrid configurations, with regional control available for self-hosted deployments. The administrative GUI mirrors the REST API, and policy management can be automated through Terraform or code-driven workflows. Policy-as-code is mature, SDKs are available in major programming languages, and support and documentation are provided in English.

The near-term roadmap focuses on productized agentic AI capabilities, including general availability of MCP and LLM Gateway products, expanded ISV customer-connect use cases, OT-specific capabilities including broader Layer 2 support, and further API and telemetry enhancements.

NetFoundry is best suited to organizations whose Zero Trust requirements center on workload connectivity, machine identity, east-west segmentation, and API security across heterogeneous and distributed infrastructure, including OT, IoT, multi-cloud, and partner domains. It is particularly relevant for ISVs delivering services into customer environments, organizations with OT/IT convergence requirements, and enterprises building agentic AI infrastructure.


Strengths

- Private keys are generated on the originating endpoint and never leave it, meaning NetFoundry cannot decrypt customer traffic
- Service-level microsegmentation enforced natively for workload-to-workload, service-to-service, API-to-API, and machine-to-machine flows from a single overlay fabric
- MCP Gateway and LLM Gateway make LLM endpoints and MCP servers dark by default, issuing strong cryptographic identities to AI agents
- Native OT and IoT support via NetFoundry zLAN
- SDK availability enables embedding of Zero Trust enforcement directly into applications
- ISV customer-connect use case support allows NetFoundry to be embedded in third-party products

Challenges

- Behavioral analytics is not available in the platform, relying on external SIEM and SOAR
- Workload, API, OT/IoT, and developer-embedded use cases are the primary focus; human remote access is a secondary capability
- Native DLP, malware scanning, and content inspection are not provided, relying on external tools where inline inspection is required
- Limited market presence and brand recognition; primarily known in OT/IoT and developer communities

Netskope – Netskope One Platform



Founded in 2012 and headquartered in Santa Clara, California, Netskope is a global cybersecurity company with the portfolio consolidated under the Netskope One brand. The platform converges SSE, SD-WAN, data security, and AI security capabilities into a unified architecture. The SSE layer includes Next-Generation SWG, CASB, ZTNA through Netskope One Private Access, FWaaS, DNS Security, RBI, Enterprise Browser, DLP, DSPM, UEBA, and cloud threat protection. The SD-WAN layer integrates through Netskope One Gateway and the unified Netskope One Client and also hosts the Device Intelligence solution. Netskope One AI Security adds dedicated capabilities for AI governance and security. Netskope One Advanced Analytics and Digital Experience Management (DEM) offer comprehensive visibility and insight across the platform.

Netskope One is built on NewEdge, the company's private security cloud, with a broad global PoP footprint, extensive network adjacencies, and direct peering with major cloud providers and internet exchanges. This gives it the performance and resilience advantages over architectures that rely heavily on public internet. NewEdge AI Fast Path optimizes latency to AI destinations. Private application access is delivered through a broker model in which Netskope Publishers are deployed as lightweight connectors (VM or container) inside customer environments and establish outbound-only TLS tunnels to the nearest NewEdge PoP. Protected applications remain invisible to unauthenticated parties, with the Zero Trust Engine evaluating access before connectivity is established.

The Netskope One Client is an agent for Windows, macOS, Android, iOS, Linux, and ChromeOS, covering SWG, CASB, Private Access, FWaaS, DNS Security, endpoint DLP, and SD-WAN. Agentless browser-based access supports web applications and non-web protocols, including RDP, SSH, Telnet, and VNC. A local broker capability delivers Universal ZTNA for on-premises users, eliminating cloud hairpinning for local traffic. Netskope One Gateway is available in hardware and virtual form factors for branches, data centers, and cloud deployments. API connectors are available for out-of-band visibility into managed cloud applications such as Microsoft 365, Box, Salesforce, Google Workspace, Anthropic Claude, and AWS.

The Zero Trust Engine is Netskope's central PDP. It evaluates access in real time using user identity, User Confidence Index scores from UEBA, device posture from MDM and EDR integrations, application risk from the Cloud Confidence Index, instance awareness distinguishing corporate from personal application instances, activity-level controls, data sensitivity from inline DLP, network location, and threat intelligence. Policy granularity extends from domain and URL level down to specific application activities, with controls differentiated by user group, device type and posture, application instance, and data sensitivity. Continuous session re-evaluation operates through the User Confidence Index and Cloud Risk Exchange, allowing external EDR, IAM, and device management risk signals to influence access decisions in near real time. Step-up authentication can be triggered mid-session, while real-time user coaching provides a softer alternative to blocking for medium-risk actions. JIT access and approval workflows are supported through ITSM integration via Cloud Ticket Orchestrator.

The User Confidence Index uses behavioral analytics to produce per-user risk scores based on data activity anomalies, insider risk indicators, credential abuse signals, and behavioral baselines. Netskope One Device Intelligence provides agentless discovery, classification, and risk assessment for IT, IoT, and OT devices at branches or campuses, with integration into SD-LAN segmentation policies enforced through Netskope One Gateway or multi-vendor network infrastructure.

Segmentation operates across user-to-application, network, and device layers. Netskope One Private Access enforces application-level access, connecting users only to explicitly authorized private applications and limiting lateral movement. Netskope One Gateway supports VRF-based segmentation across endpoints, branches, data centers, and clouds, enabling segment-aware topologies and dynamic tunnel policies. IoT and OT microsegmentation is automated through Device Intelligence, which discovers and

fingerprints unmanaged devices through passive network observation (such as cameras, PLCs, and medical equipment that cannot host a Netskope client) and assigns each a dynamic risk score that drives SD-LAN policy enforcement. Zero Trust SASE SIM extends segmentation and enforcement to IoT and OT devices over cellular networks.

Monitoring and analytics are provided through Netskope One Advanced Analytics, which offers cloud risk posture visibility, behavioral intelligence, custom reporting, and forensics for data movement analysis. AgentSkope is a framework for introducing agentic automation for security and network operations, including agents for DLP workflows, insider threat investigation, CCI insights, private access operations, and DEM insights, with more to come. Cloud Exchange enables bidirectional threat intelligence sharing, ITSM workflow integration, risk score exchange, and near-real-time log export to SIEM, XDR, and analytics platforms. Automated responses include session termination, step-up authentication, user coaching, account blocking, and playbook-based incident response.

Threat prevention is delivered through a multi-layered inline stack covering malware detection, multistage sandboxing, phishing detection, IDS/IPS, DNS security, RBI, and FWaaS. Netskope's strength is the consistent application of inspection and enforcement across web, SaaS, private applications, cloud traffic, and branch flows. FWaaS supports user-based access controls with application identification on standard and non-standard ports. Cloud TAP supports packet capture and session key export for compliance and forensic use cases. Private application traffic can receive the same inline inspection used for web and SaaS destinations.

A single DLP policy engine covers web inline, SaaS at rest and inline, IaaS and PaaS, private applications, email, and endpoints under unified policy and incident management. The classification engine supports content types, data identifiers, compliance templates, exact data matching, fingerprinting, Optical Character Recognition (OCR), and ML-based classifiers. File-level encryption and tokenization for structured fields are available. DLP On Demand allows customers to integrate Netskope DLP into applications and on-premises environments without routing traffic through the Netskope cloud. Visibility and control over data used to train public and private LLMs are offered, including protection against data poisoning. Browser session controls include copy-paste, download, upload, printing, and screenshot prevention through the session proxy and Enterprise Browser. Access policies can reference sensitivity labels from Microsoft Purview, Fortra Titus, Box Shield, and Google Labels.

Agentic Broker provides visibility and control over MCP transactions, helping organizations govern how autonomous AI systems access enterprise data and tools. AI Guardrails moderate user and agent interactions with LLMs, addressing prompt injection, jailbreaks, and inappropriate or risky outputs. AI Gateway enforces policies for privately hosted LLMs and AI applications, centralizing authentication and traffic management for agentic communications in on-premises or VPC-hosted environments. AI Red Teaming simulates adversarial attacks, including multi-turn jailbreaks, to identify weaknesses before production. Shadow AI governance provides visibility into AI applications, inline DLP for prompts and uploads, and blocking of high-risk AI destinations.

Netskope holds a broad certification portfolio including ISO 27001, ISO 27017, ISO 27018, SOC 2 Type II, FedRAMP High, Germany C5, UK Cyber Essentials, CSA STAR Level 2, IRAP, Spain ENS, and mappings to GDPR, DORA, the EU AI Act, and other regional frameworks. The SSE platform is delivered as a cloud service, while Netskope One Gateway and DLP On Demand support on-premises and private cloud use cases. Management planes are available in multiple countries for data residency. Administration is unified through a single console and policy engine across SSE, SD-WAN, data security, and AI security, with REST APIs, MCP Server, GraphQL, and Cloud Exchange integrations. Support is global, with multilingual documentation.

Netskope One is best suited to large and mid-market enterprises seeking a cloud-native SSE or SASE platform with deep application intelligence, granular data protection, and rapidly maturing AI security controls. The Zero Trust Engine, DLP breadth, Data Lineage, NewEdge network, and AI Security suite make it especially relevant for organizations with complex data governance, insider risk, shadow AI, or agentic AI adoption challenges.


Strengths

- Private security cloud with broad global PoP footprint and direct peering with major cloud providers provides lower latency and high resilience
- DLP On Demand enables integration of Netskope DLP directly into customer applications and environments without routing traffic through the Netskope cloud
- Zero Trust SASE SIM extends segmentation and enforcement to IoT and OT devices over cellular networks
- Broad LLM and Agentic AI governance and security capabilities
- AgentSkope introduces autonomous agentic workflows for DLP investigation, insider threat analysis, CCI insights, and private access operations
- Device Intelligence (HyperContext, TruID) extends Zero Trust to unmanaged IT, IoT, and OT devices through agentless network-based fingerprinting and dynamic risk scoring
- Strong NHI governance, including AI agents, MCP, and LLM gateways
- Broad certification and assurance portfolio with data residency options across multiple regions

Challenges

- The platform's depth and breadth can create operational complexity when compared to less sophisticated offerings
- Data masking, redaction, and watermarking are not natively available, requiring third-party tools
- East-west workload-to-workload segmentation inside data centers and cloud environments is not natively supported, relying instead on integration with alliance partners
- Advanced device posture and endpoint risk signals depend on third-party MDM and EDR integrations rather than native endpoint trust capabilities

Pomerium – Pomerium



Leadership



Founded in 2019 and headquartered in Beaverton, Oregon, Pomerium is a privately held company built around an open-source, identity-aware access platform. The commercial offering extends the open-source core with enterprise management, hosted control plane options, and advanced policy capabilities. The customer base is concentrated in North America, with some traction in EMEA and APAC, and spans financial services, aerospace and defense, government, pharmaceutical, healthcare, and software development industries.

Pomerium's architecture separates the control plane, where policies are defined and managed, from the data plane, which is always self-hosted by the customer and deployed in front of protected resources. The control plane can be self-hosted or consumed as

Pomerium Zero, a hosted service in which the customer still retains the data plane in their own infrastructure. Policies are cached in the data plane, so enforcement continues if the control plane becomes temporarily unavailable. TLS is terminated at the proxy and re-encrypted, with encryption keys held by the customer. Pomerium does not receive application payload data. Enforcement models include proxy-based, agent-based access through a CLI tunnel for protocols like RDP, MySQL, Redis, and UDP, sidecar deployment for containerized workloads, and Kubernetes-native deployment, including protection of the Kubernetes API and kubectl.

Access policy is authored in Pomerium Policy Language, a YAML-based declarative syntax designed to be readable by both security teams and developers, or in Rego for organizations using Open Policy Agent. Policies are organized in hierarchical namespaces, allowing global security teams to define baseline controls that cascade to application-level policies managed by individual teams, without allowing lower-level policies to override global rules. Policy granularity extends from host and application level down to HTTP paths and methods, specific MCP tool calls within agentic sessions, and specific SSH commands.

Authorization is evaluated for each request rather than per session, which is a key distinction from network-level access tools. JIT and time-bound access are supported, as well as instant revocation through API and approval workflows. Policy changes are versioned and auditable, with change history available for compliance. Policies can be managed through Terraform, Helm charts, Kubernetes Ingress controllers, and GitOps workflows.

Device posture signals can be consumed from MDM and device management tools, including a built-in FleetDM integration. EDR and XDR signals, identity risk signals, HR data such as out-of-office status, and custom JSON-format signals from external sources can be incorporated into policy decisions through a plugin architecture. When posture checks fail, Pomerium can display a branded remediation page explaining which policy condition failed and how the user can restore compliance. Posture can also dynamically restrict functionality at the sub-request level, including SSH command filtering and Kubernetes API access.

Because the proxy sits directly in front of each protected application or service, each resource effectively becomes a segment of one. Users have no visibility into unauthorized resources, regardless of network location. Workload-to-workload and service-to-service communication is supported through machine-to-machine access using OIDC service accounts and JWT-based authentication, with the same per-request policy evaluation applied. Kubernetes environments are a primary use case, with the Kubernetes Ingress controller enforcing policies at the cluster level. East-west network traffic control is explicitly out of scope.

Session logs, policy decision logs, and administrative activity logs are captured, and can be streamed in JSON format to log collectors or SIEM platforms. OpenTelemetry tracing supports distributed request tracing across Pomerium components, and customers can connect existing observability tools to Pomerium metrics endpoints. The platform does not provide anomaly detection or full analytics beyond real-time session visibility. In practice, organizations commonly connect Pomerium to Grafana, Splunk, Elastic, or similar tools for trend analysis, investigation, and reporting. Automated response actions include session

termination through policy re-evaluation, identity blocking through explicit deny policies, and device isolation.

Since the proxy terminates TLS from the client and re-encrypts traffic upstream, it can inspect request metadata, headers, and protocol-level signals before making authorization decisions. Continuous re-evaluation reduces the usefulness of stolen credentials or session tokens once identity or posture conditions change. Sensitive authentication cookies are stripped before requests are forwarded to upstream applications, limiting token exposure. For MCP and agentic traffic, Pomerium inspects the protocol structure of MCP requests inline and can enforce policy on specific tool calls based on the agent's verified identity, including blocking individual MCP tools.

Pomerium does not perform DLP, content masking, or redaction. Policies can reference data classification labels or sensitivity metadata where those labels are available as structured request context, enabling label-aware routing decisions.

NHI protection is a first-class capability and an area of active investment. Workloads, APIs, service accounts, and AI agents are treated as identities subject to the same per-request policy enforcement as human users. NHIs are issued through the enterprise console or management API and authenticate using JWTs. Workloads running in AWS, Azure, and GCP can federate using OIDC-compliant tokens from the cloud provider's metadata service, enabling secretless cross-cloud access without static API keys; other NHI types, including on-prem services, containers, and AI agents, authenticate through the JWT-based service account model described above. For agentic AI, Pomerium can act as an MCP gateway, enforcing per-tool-call authorization based on the agent's verified identity and logging all agent actions at the call level. A hosted cluster capability for MCP deployments is currently in beta, while PAM-style session recording and SSH command restriction capabilities are under development.

Pomerium holds SOC 2 Type I and Type II attestation across all five trust service criteria. Because the data plane is always customer-hosted and Pomerium does not receive application payload data, customers retain full control over data residency regardless of where the control plane runs. Support is available in English, French, Italian, and Russian, while documentation is in English. On-site support is available globally at additional cost. The administrative experience is developer-native, centered on policy-as-code, REST APIs, Terraform providers, Helm charts, SDKs, and GitOps compatibility. The commercial Enterprise and Zero editions include a full-featured admin GUI; the open-source Pomerium Core remains configuration-driven. The recently added Routes Portal additionally gives end users a browsable catalog of authorized services.

The roadmap focuses on expanding PAM-style capabilities, including session recording, SSH command restrictions, and deeper protocol-level controls, alongside continued development of the agentic gateway, MCP-specific tunneling, broader MCP tool policy criteria, and transition of the hosted MCP cluster capability out of beta.

Pomerium is best suited to organizations seeking identity-aware, per-request access enforcement at the application layer without deploying a client on end-user devices or

modifying protected servers. It is particularly relevant for complex multi-cloud, hybrid, and legacy application environments, high data sovereignty or air-gap requirements, Kubernetes-heavy infrastructure, and agentic AI systems requiring granular authorization. Its open-source nature and developer-native policy model make it a strong fit for teams operating in code-driven environments.


Strengths

- Every individual HTTP request, SSH command, or tool invocation is independently authorized against current policies
- Pomerium never holds customer keys or receives payload data, regardless of whether the control plane is self-hosted or vendor-managed
- Native MCP gateway capability enforces per-tool-call authorization for AI agents based on verified agent identity, with full call-level audit logging
- Policy-as-code using Pomerium Policy Language (YAML) or Rego, with hierarchical namespaces to allow delegating policy authoring to application teams
- Developer-native operational model with Terraform, Kubernetes, GitOps, and SDK support makes it particularly well suited to code-driven infrastructure teams

Challenges

- A full-featured admin GUI is available only in the commercial Enterprise and Zero editions; the open-source Core remains code-configured
- No DLP, content inspection, masking, or redaction; the platform makes label-aware routing decisions but does not process content
- Lacks east-west network traffic controls
- UEBA and anomaly detection require streaming to external tooling; built-in dashboards cover traffic, runtime, and session telemetry
- Limited market presence outside North America; primarily known in developer-centric and security engineering communities

Sophos – Sophos Workspace Protection (ZTNA)



Leadership



Founded in 1985 and headquartered in Abingdon, United Kingdom, Sophos is a privately held cybersecurity company with a long-standing focus on endpoint, network, email, cloud, and managed security services. The company operates through a large global channel ecosystem of MSPs, resellers, and system integrators across North America, EMEA, APAC, and Latin America. Its sweet spot remains small and mid-sized organizations, although Sophos also serves larger enterprises.

The most significant recent development in Sophos's Zero Trust access architecture is Sophos Workspace Protection, a bundled product suite that consolidates Sophos ZTNA, Sophos Protected Browser, Sophos DNS Protection for Endpoints, and Sophos Email Monitoring System under a single license. The product is channel and MSP-delivered and is managed through Sophos Central, the company's unified cloud management platform, hosted on AWS with regional data center options for data residency.

The platform is delivered as a cloud-brokered service. The ZTNA gateway can be deployed as a virtual appliance on VMware or Hyper-V, as a cloud-hosted instance on AWS, or natively integrated into Sophos Firewall v20 and later. This last option eliminates the need for a separate gateway component in environments already using Sophos network security appliances. Gateways establish outbound-only connections to the Sophos cloud, making protected applications invisible to the public internet. Organizations already using Sophos endpoint protection can enable ZTNA as an additional licensed capability without deploying a separate client.

Agentless access is provided through Sophos Protected Browser, a hardened Chromium-based enterprise browser built in partnership with Island.io. It provides ZTNA-integrated access to private web applications, RDP, and SSH, enabling BYOD, contractor, and unmanaged device usage. Identity integration uses OIDC-compatible providers, including Microsoft Entra ID and Okta, with MFA required for access.

Policy enforcement combines user group membership with device posture signals from the endpoint agent and, in browser-based access scenarios, the Sophos Protected Browser. A notable architectural differentiator is Synchronized Security Heartbeat, through which Sophos endpoint agents continuously report health status to ZTNA, Protected Browser, and Sophos Firewall. If an endpoint's health degrades during a session because of a detected threat, ransomware behavior, or policy violation, ZTNA can revoke application access automatically. Device posture checks can evaluate endpoint protection presence and health across Sophos and supported third-party endpoint tools, disk encryption status, and OS-level conditions. Posture requirements can vary per application. These posture requirements can be enforced on a per-application basis across both ZTNA and Protected Browser access paths.

Segmentation and lateral movement control operate across two complementary layers. Within ZTNA, application-level policies prevent users from reaching anything beyond their permitted applications. Within the local network, Synchronized Security extends automated isolation to the firewall and switching layer. When the endpoint agent detects ransomware or active adversary behavior, Sophos Firewall can isolate the affected device at the network level and share its MAC address with other Synchronized Security-enabled endpoints, which stop Layer 2 communication with the compromised host. Active Threat Response allows MDR analysts and third-party SOC teams to push threat feeds of malicious IPs, URLs, and domains directly to Sophos Firewall through API integration. NDR Essentials, available on XGS appliances, provides encrypted payload analysis by offloading TLS and DNS traffic analysis to the Sophos cloud without inline decryption.

Sophos Protected Browser is the browser-centric enforcement component of Workspace Protection. Integrated with Sophos threat intelligence, it provides access to private web applications, RDP, and SSH, while adding SaaS application usage monitoring and Data Boundary Controls. These controls govern information flows between defined application groups, covering copy-paste, file upload and download, screen capture, printing, and page saving without requiring full inline content inspection. Administrators can define where corporate data can and cannot move, including policies for GenAI tools. Protected Browser also provides web and DNS filtering and integrates with Sophos XDR and MDR for unified visibility.

Sophos Endpoint, including Intercept X, provides CryptoGuard ransomware protection with rollback capability, anti-exploitation controls, Adaptive Attack Protection for active hands-on-

keyboard behavior, and Critical Attack Warning for attacks spanning multiple systems. Sophos Firewall adds TLS inspection, IPS, web filtering, DNS security, AI-powered sandboxing through SophosLabs Intelix, and application identification through a single-pass DPI engine. Active Threat Response capabilities allow real-time threat feeds from Sophos MDR analysts to be pushed to customer firewalls. Third-party threat feeds from sector-specific and regional sources are also supported.

Monitoring and analytics draw on the Sophos Central data lake, which aggregates telemetry from Sophos products and supported third-party integrations. Cases and investigations are presented as collaborative workspaces with threat graphs, supported by automated root cause analysis and confidence-level attribution. Agentic AI capabilities include autonomous investigation workflows, continuous agentic threat hunting, and a GenAI quality assurance agent that reviews human-led MDR investigations for consistency. Sophos ITDR extends visibility into identity risk, compromised credential exposure, and Microsoft 365 account containment.

Sophos Endpoint Protection includes peripheral control (such as removable media) and web-based data transfer controls; together with the Data Boundary Controls in Protected Browser, these provide channel-level DLP that monitors and restricts where sensitive data can move. Payload-level capabilities such as full content inspection, masking, redaction, watermarking, and classification-driven enforcement are not provided within the ZTNA or browser layers. NHI governance, identity-aware workload microsegmentation beyond endpoint-integrated ZTNA, and dedicated agentic AI authorization controls are not present in the Sophos Zero Trust portfolio today. The Secureworks acquisition expands the broader detection and response layer but does not close these gaps. However, Sophos NDR helps address related operational use cases through east-west traffic visibility, lateral-movement detection, and coordinated response actions, and AI governance across Protected Browser, Endpoint, and Firewall provides allow/warn/deny policy and browser-layer data boundary controls for AI application use.

Sophos holds ISO 27001, SOC 2 Type II, PCI DSS, and Germany C5 certifications. Support is available 24/7 in multiple languages, and documentation is localized across several major languages. The roadmap reflects the Secureworks integration, including expanded ITDR coverage, broader Microsoft 365 identity response, enhanced threat hunting, broader third-party integrations, expanded automated response actions across additional control planes, and continued development of GenAI and agentic security operations.

Sophos Workspace Protection provides a pragmatic, browser-centric alternative to SASE complexity. It is best suited to small and mid-sized organizations, and to MSPs or channel partners serving them, that want a tightly integrated cybersecurity defense system that coordinates automated responses across ZTNA, endpoint protection, firewall, browser-based access control, and 24/7 MDR, all via a single management interface.

Strengths


- Native ZTNA gateway integration into Sophos Firewall eliminates the need for a separate deployment in environments already using Sophos network security
- Synchronized Security Heartbeat provides real-time endpoint health telemetry, enabling automatic access revocation when attacks are detected mid-session

- Automated network-level isolation at the firewall and switching layer triggered directly by endpoint threat detection
- Sophos Protected Browser provides ZTNA-integrated browser access for BYOD and unmanaged devices
- CryptoGuard ransomware protection with rollback capability and Adaptive Attack Protection directly complement ZTNA enforcement
- Channel and MSP-native delivery model makes the platform highly accessible to smaller organizations through managed services without requiring enterprise IT resources

Challenges

- No NHI governance, workload microsegmentation beyond endpoint-integrated ZTNA, or agentic AI authorization controls
- No inline content inspection, masking, or redaction in the ZTNA or browser layers; data controls operate at channel and boundary level only
- Continuous mid-session risk evaluation via Synchronized Security Heartbeat requires the Sophos endpoint agent
- East-west workload segmentation is handled via Synchronized Security at the endpoint and firewall layer, not through dedicated identity-aware microsegmentation
- Primarily strong in the SMB and mid-market segment; large enterprise customers with complex hybrid infrastructure may find the platform's scope limiting

Specops – Specops Device Trust



Specops is the identity and access management division of Outpost24, a Swedish cybersecurity company established in 2001 and headquartered in Stockholm. The product evaluated here is Infinipoint, which is being rebranded and integrated as Specops Device Trust following its acquisition by Specops/Outpost24. This acquisition expands Specops beyond its established identity security and password management portfolio into device identity, posture verification, and Zero Trust access enforcement. Specops serves customers globally, with a strong presence in North America and EMEA and smaller coverage in APAC and Latin America.

Infinipoint is a cloud-native, agent-based Zero Trust access platform positioned as a device trust broker. It sits between users and applications at authentication time, performing real-time device identity verification, posture assessment, and guided remediation before access is granted. The platform is delivered as a SaaS service on GCP, with logical tenant isolation and physical separation between US and EU regions for data residency. The enforcement model relies on a lightweight agent for Windows, macOS, Linux, iOS, and Android, covering both managed corporate devices and unmanaged or BYOD endpoints.

The agent operates in two modes. Corporate mode can execute remediation scripts and collect broader telemetry, while privacy-first BYOD mode limits data collection and performs guided rather than automatic remediation. The agent cannot wipe devices in BYOD mode; a design choice communicated to users through a privacy disclosure page. Agentless or browser-only enforcement is not supported, since browser sandboxing limits prevent OS-level posture checks without an installed component. The architecture separates edge enforcement gateways, a centralized control plane for policy and identity lifecycle, and a Kafka-based data plane with multi-tier telemetry storage. Control plane resilience is provided through Kubernetes multi-AZ deployment and automated failover.

The core premise of Infinipoint is binding user identity to specific trusted devices. This is intended to reduce the risk of credential phishing, session token theft, and access from unmanaged or unknown devices. Device-user pinning means that a stolen password, MFA token, or session cookie cannot simply be replayed from an attacker's machine. Administrators can define per-user or per-group device quotas and registration policies, including optional manual approval for new device registrations. Specific compliance items can also be exempted for individual users, devices, or groups without modifying the underlying ruleset or granting a global bypass. Users who exceed their quota or attempt access from an unregistered device can be blocked and directed to an administrator or to a self-service device management portal.

Each access policy rule combines a configuration, defining posture checks or authentication requirements, with a target scope, specifying users, groups, or device types. If no rule matches, requests fall through to a configurable default, which can be set to deny. Policy types include user authentication, device authentication, ownership and pinning, and device posture for workstations and mobile devices. Posture policies can vary by application, allowing permissive checks for low-risk applications and stricter requirements for sensitive resources.

JIT access is supported for user-to-application scenarios. Continuous session re-evaluation runs at configurable intervals (10 minutes by default), with posture drift signals sent to integrated ZTNA or SASE solutions such as Cloudflare, Netskope, and Zscaler that can trigger re-authentication or access revocation. Step-up authentication is supported at login but not mid-session in response to posture changes. Approval workflows for sensitive access requests are not currently available, though exception approval is on the roadmap.

Infinipoint provides a broad library of built-in compliance checks across Windows, macOS, Linux, iOS, and Android. These cover OS version and patch status, disk encryption, firewall and security service state, software version checks, CVE-based vulnerability assessment, jailbreak detection for mobile devices, registry entries, and file existence. CIS Benchmark coverage is also included for assessing endpoint configurations against industry-standard baselines. EDR integrations include CrowdStrike, SentinelOne, and Cortex XDR. MDM and UEM integrations include Microsoft Intune, Jamf, and Kandji.

When posture failures occur, the platform presents users with a remediation screen and, where supported, a "Fix Now" button that applies the correction automatically. Administrators can define grace periods, allowing users temporary access despite outstanding issues. This

model is key to the platform's value proposition: it lets security teams enforce posture requirements without immediately overwhelming service desks or blocking business operations.

Administration is performed through a web console that displays device inventory, user management, vulnerability dashboards, policy configuration, compliance reporting, and access and audit logs. Each access attempt records the relevant identity, device state, policy evaluation result, and reason for any denial. Data is retained for 18 months. SIEM integration is available through webhooks, and REST APIs support log export and policy management. GenAI capabilities in the console assist with policy authoring and configuration.

Session hijacking and replay attacks are mitigated by requiring authentication to originate from a verified device associated with the user. Data-centric controls are not present. Lateral movement is addressed indirectly: an attacker with credentials cannot use them from a different device to access additional applications.

NHI support is limited. Service accounts can be grouped and governed through policies, and machine actions are logged for audit purposes. Workload identities, API identities, and AI agent governance are not currently supported, although the vendor identifies these areas as roadmap priorities.

Identity provider integrations include Okta, Microsoft Entra ID, Ping Identity, OneLogin, and Google Workspace Cloud Identity, with bidirectional integration available for Okta and Entra. Integrations with ZTNA and SASE platforms such as Cloudflare, Netskope, and Zscaler allow Infinipoint to trigger access revocation or re-authentication when posture changes are detected during a session. ServiceNow integration enables users to raise support tickets directly from the access interface with posture context included automatically. Policy management is available through REST APIs.

The platform holds SOC 2 Type II certification across all five trust service criteria. Keys are managed through Google Cloud KMS with hardware-based key protection on Google infrastructure. Support is available in English, German, French, Spanish, Vietnamese, and Hebrew, while documentation is available in English, German, and French. Standard support is included, with a premium 24/7 support option available.

The near-term roadmap includes multilingual end-user interface support, desktop login integration that embeds posture checks into OS authentication, agent support for non-privileged users with fallback to manual remediation when admin privileges are required, enhanced controls for service accounts and AI agents, and silent auto-remediation. Medium-term work includes tighter integration with Specops authentication platform and service desk verification capabilities.

Specops Device Trust is best suited to organizations seeking to close the device trust gap between their identity provider and MDM platform, especially in heterogeneous environments spanning managed devices, BYOD endpoints, contractors, and personal mobile devices. Its strongest differentiator is consistent cross-platform posture enforcement, including Linux support, combined with a BYOD-friendly privacy model. Specops is best

viewed as a device trust and posture layer that strengthens existing identity and access architectures rather than replacing them.


Strengths

- Device-user binding architecture means stolen credentials, MFA tokens, or session cookies cannot be replayed from an attacker's machine
- Privacy model with distinct Corporate and BYOD agent modes limits data collection on personal devices without compromising enforcement
- Guided self-remediation with "Fix Now" automation and grace periods allows policy enforcement without productivity disruption
- Consistent cross-platform posture enforcement, including Linux support
- EDR and MDM integrations with most popular vendors enable posture checks against managed security tool state
- Continuous posture drift signals can be pushed to third-party ZTNA and SASE platforms (Cloudflare, Netskope, Zscaler) to trigger session revocation or re-authentication

Challenges

- Works as a device trust and posture broker that must be paired with a separate access platform
- No agentless or browser-based access path for unmanaged devices; the platform requires an installed agent
- Step-up authentication applies at login but not mid-session
- Data-centric controls (content inspection, DLP, masking) are not part of the platform
- Actual session revocation on posture drift depends on an integrated third-party ZTNA or SASE

Systancia – cyberelements



Founded in 1998 and headquartered in Pfaffstätt, France, Systancia is an independent software vendor focused on European sovereignty, secure access, and identity-centric security. Its cyberelements platform is positioned as a European-native Zero Trust platform, combining PAM, Identity Governance and Administration (IGA), Single Sign-On (SSO), and Zero Trust access enforcement in a single product. Systancia's customer base is concentrated in EMEA, anchored in France with expanding presence across Europe, including recent wins in Italy and Spain, and traction across regulated and industrial sectors.

The cyberelements architecture is built on what Systancia calls a double barrier design, separating authentication and authorization from resource connectivity across two control layers. The first layer is the Mediation Controller, a cloud-hosted or self-hosted control plane that manages identity verification, policy evaluation, session orchestration, multi-tenancy, and audit. The second layer is the Edge Gateway, a hardened virtual appliance deployed at customer sites, in cloud environments, or at network perimeters in front of protected resources. Edge Gateways connect outbound only to the Mediation Controller on port 443 and require no inbound exposure. Access to a resource is initiated only by authenticated and authorized users through the controller, which then establishes a tunneled connection

through the appropriate gateway. Users and attackers cannot connect directly to a gateway or to resources behind it.

The architecture is multi-tenant, with organizational and network segmentation managed through distinct controller tenants and gateway configurations. This enables a single platform instance to isolate IT from OT, regulated from non-regulated environments, or multiple customer organizations from one another. The Mediation Controller can be consumed as SaaS with rapid tenant provisioning, deployed as self-managed software on any cloud or on-premises infrastructure, or used in a hybrid deployment.

Enforcement models include agent-based, agentless clientless browser-based, proxy-based, sidecar, and network-integrated approaches. Clientless access uses an HTML5 gateway with protocol break: the connection from the user endpoint terminates at the controller layer, and a new connection is initiated from the gateway to the protected resource. This prevents payloads from unmanaged or contractor devices from reaching target systems directly. Full protocol client access is supported for RDP, SSH, VNC, and SMB. A generic tunneling capability extends enforcement to database protocols as well as proprietary and legacy protocols, including OT protocols, through encrypted end-to-end tunnels with port forwarding. Data in transit is protected by TLS 1.3. Volatile and random destination ports are used for client-side connections, reducing the viability of brute force attacks against exposed services.

Access policy combines RBAC, ABAC, contextual conditional access, and risk-based evaluation. Policy granularity extends from server and system level to application and API route level. JIT access is supported through dynamic access policies and workflow approval chains. Session re-evaluation occurs throughout active sessions, with triggers including identity risk changes, device posture changes, network or location changes, and behavioral anomalies. Step-up authentication and approval workflows for sensitive access requests are supported. Privileged access scenarios inherit the same Zero Trust infrastructure characteristics, including protocol break, encrypted tunneling, clientless access, URL rewriting, and volatile ports, reducing the need for traditional jump servers.

Context and posture evaluation is performed natively without separate MDM, UEM, or EDR tools, although external signals from such platforms can be incorporated. Posture is evaluated at authentication and continuously during the session at configurable intervals, with re-evaluation triggered server-side rather than on the client, which improves tamper resistance. Different requirements can be defined for each resource, or when posture conditions fail during a session, cyberelements can revoke access, trigger step-up authentication, or generate alerts. Managed, unmanaged, BYOD, embedded, and IoT devices are supported.

A network filtering agent adds endpoint-level segmentation, allowing individual users within the same group to receive different lateral access permissions. East-west access control within the platform's scope is supported through the tunneling and gateway architecture for workload-to-workload communication. Kubernetes-native microsegmentation is out of scope by design, reflecting an enforcement focus on regulated IT/OT and legacy or proprietary

protocols. Application dependency mapping and policy recommendation are on the roadmap.

Monitoring and analytics cover session logs, flow logs, policy decision records, and identity context for all access activity. Real-time session visibility is available through the administrative console. Behavioral baselining and anomaly detection use ML models applied to session metadata and access patterns. Session recording is available for privileged access, and the Neomia Lens module applies AI video analysis to recorded sessions in real time. It evaluates session activity against the defined mission or task the user was authorized to perform, as well as security best practices including MITRE ATT&CK mappings. Telemetry can be exported through syslog, APIs, and log streaming to SIEM platforms. Automated responses include session termination, step-up authentication, account or identity blocking, device isolation, and AI-driven triage of session alerts through the Neomia Lens module.

Brute force attacks are mitigated by assigning volatile, randomized destination ports to client-side connections so that listening services cannot be enumerated. CSRF, XSS, and cookie theft risks are reduced through URL rewriting that removes predictable resource paths and through session tokens scoped to the authenticated TLS context. Supply chain and malware exposure from unmanaged endpoints is reduced because the HTML5 clientless gateway terminates the user session at the controller layer and initiates a separate connection to the target resource (protocol break). Neomia Pulse, the platform's deviceless biometric MFA module, authenticates users through behavioral biometric profiling of keyboard dynamics combined with contextual signals such as IP address and location. Inline payload inspection is not native; it is delivered through third-party integrations over ICAP.

Data in transit is protected by TLS 1.3 and optional customer-managed keys. Browser session controls for copy-paste and upload or download can be provided through embedded Virtual Desktop Infrastructure isolation or RBI, but screenshot prevention and print controls are not native. Native DLP, masking, redaction, and watermarking are not provided; these are delivered through partner integrations, while the double-barrier protocol break combined with RBI/VDI mitigates the unmanaged-endpoint payload threat by design. Workloads, APIs, service accounts, and AI agents can be governed through the same Zero Trust access framework as human users. The MCP Proxy is an enforcement mechanism for agentic AI, placing one proxy instance on the user's desktop alongside the AI client and a second in front of the MCP server connectors. Agent-to-tool communications are routed through the proxy and double barrier architecture, producing an audit trail of agent actions, tool calls, and connector interactions. Workload identity federation is supported, and least-privilege enforcement for machine and agent actions is applied through the access policy engine. Intent-based access control for agent actions is planned.

Systancia holds SOC 2 Type II attestation, the France Cybersecurity Label, and SecNumCloud compliance for its SaaS hosted infrastructure. Both ISO 27001 and FITCEM-level product certification (French ANSSI CSPN) are under way. Support is available in English, German, and French, with business-hours and 24/7 support options. Documentation is available in English and French, and on-site professional services are available within EMEA. The roadmap focuses on geo-distributed Mediation Controller points of presence, AI-

driven role and entitlement analysis, ephemeral account provisioning combining IGA and PAM, Windows application SSO by injection, and expansion of the agentic AI MCP connector framework.

cyberelements is best suited to regulated and sovereignty-driven organizations across mid-market and large enterprises, especially those with French or northern European regulatory requirements, mixed IT and OT environments, or strong sovereignty needs. Its strengths are native multi-tenancy, MSSP-friendly isolation, a genuine Zero Trust access architecture, and the combination of PAM, IGA, SSO, and access enforcement in one platform.


Strengths

- Double-barrier architecture with protocol break at the controller layer, an isolation model stronger than most proxy-based ZTNA approaches
- Native PAM, IGA, SSO, and Zero Trust access enforcement in a single product eliminates the need to integrate multiple identity and access control
- MCP Proxy capability places PEPs at both ends of the AI chain, creating a double-sided audit trail
- Real-time AI video analysis of privileged sessions, evaluating actions against the defined mission and configurable security framework mappings
- Deviceless biometric MFA enables strong authentication without hardware requirements
- Addresses all French sovereign cloud requirements

Challenges

- Platform is primarily known in France and across Europe; limited market presence beyond that
- No native inline payload inspection; integration available via third-party partnerships through ICAP
- DLP, masking, and redaction delivered through partner integrations rather than natively
- Kubernetes-native microsegmentation is out of scope by design; enforcement focus is on regulated IT/OT and legacy or proprietary protocols
- Screenshot prevention and print controls are not native to browser-based sessions

Teleport – Teleport Infrastructure Identity Platform



Founded in 2015 and headquartered in Oakland, California, Teleport is a privately held infrastructure identity company focused on identity-native security for engineering, platform, and security teams. Its primary market is organizations with large, dynamic, and complex infrastructure estates, including cloud-native technology companies, modern enterprises replacing legacy PAM tools, AI infrastructure providers, and organizations deploying agentic AI workflows that require governance for non-human actors.

The Teleport Infrastructure Identity Platform is built on a certificate-centric, secretless architecture. Rather than vaulting, rotating, or brokering static credentials, Teleport acts as a certificate authority that issues short-lived X.509 and SSH certificates to all identities, including human users, services, CI/CD pipelines, IoT devices, workloads, and AI agents, at the moment access is requested. These certificates encode roles and permitted actions

directly, expire automatically, and cannot be reused. This model is central to Teleport's value proposition: eliminating static credentials and standing privileges and reducing the window in which compromised identities can be abused.

The architecture consists of two core control plane components. The Auth Service maintains certificate authorities, stores cluster state, and collects audit events. The Proxy Service enables users and agents to reach protected infrastructure in private networks through reverse tunnels, avoiding inbound port exposure. Agents, deployed alongside protected resources as instances of the Teleport binary, establish outbound SSH reverse tunnels to the Proxy Service and verify client certificates against the Auth Service CAs before allowing connections.

Teleport supports multiple enforcement models: agent-based deployment for servers, databases, Kubernetes clusters, applications, and Windows desktops; agentless or browser-based access through the web UI using WebSocket Secure tunnels; identity-aware proxy mode where private keys remain on the client and connections are mutually authenticated (including Git repository access and MCP server connections); sidecar and Kubernetes DaemonSet deployment for workload identity; and network tunnel mode for resources behind firewalls. TLS routing compresses traffic to a single TLS port, supporting SSH, Kubernetes, SQL, RDP, HTTP/S, MCP, and generic TCP.

Access policy is based on RBAC with ABAC extensions and time-based controls. Policies use labels rather than network topology, making them portable across cloud, on-premises, hybrid, and air-gapped environments. Role templates can consume attributes from SAML or OIDC identity providers, allowing role assignments to follow identity provider changes automatically. Authorization granularity extends from server and system level through application and Kubernetes namespace level down to database names, tables, procedures, and views, Kubernetes resources, and specific MCP tools for AI agent access.

JIT access is implemented through Access Requests supporting temporary elevation to specific roles or resources for both users and machines. Approval workflows route through collaboration, ITSM, and incident response tools, and can support automatic approval, dual authorization, or conditional approval. Session MFA can be enforced globally or per role across SSH, Kubernetes, database, application, and desktop access. Mid-session risk signals, including identity risk, device posture changes, location changes, and behavioral anomalies, surface as alerts in Teleport Identity Security, but do not currently trigger automated session revocation.

Device Trust binds user identity to enrolled hardware through cryptographic attestation. It can be optional, required, or required only for human users, with enforcement configurable per role. Integrations with Jamf Pro and Microsoft Intune support device inventory synchronization. BYOD scenarios can be handled through tiered policies, for example, allowing unmanaged devices to reach non-production resources while requiring enrolled devices for production access.

Segmentation in Teleport is identity-aware and topology-independent. At the user-to-resource layer, label-based RBAC ensures each identity can reach only explicitly permitted

resources. At the workload-to-workload layer, Teleport Workload Identity issues SPIFFE-compliant X.509 certificates for mutual TLS between workloads, replacing shared secrets with short-lived cryptographic identities. Machine ID issues renewable certificates to non-interactive bots and automation workflows through methods such as cloud IAM attestation, TPM attestation, GitHub Actions OIDC, GitLab CI tokens, and others. Kubernetes namespace-level segmentation is enforced through the Teleport Kubernetes Service and operator. Workload federation with AWS, Google Cloud, and Azure enables short-lived cloud credentials without static API keys. Access Graph provides visual mapping of identity-to-resource paths and helps identify over-provisioned access.

Session telemetry covers SSH terminal streams, kubectl execs, SQL queries with database and user context, RDP video for Windows desktops, and HTTP request logs for application access. Real-time session visibility allows security teams to observe and join active sessions. Behavioral baselining and anomaly detection are provided, covering identity vulnerabilities, unusual access patterns, permission changes, posture failures, long-running or inactive sessions, and lateral movement indicators. AI-powered Session Summaries analyze recorded sessions and summarize commands, privilege escalations, security-relevant behaviors, and severity levels. Automated responses include identity locking, session termination, step-up authentication triggers, and device blocking.

Certificate generation counters can detect stolen renewable certificates and lock affected identities when mismatches occur. Device Trust limits credential reuse from unauthorized devices, while per-session MFA reduces certificate replay risk. Shadow access detection identifies unmanaged SSH keys, expired credentials, and unauthorized access paths. Enhanced session recording can capture process execution, file activity, and network connections during SSH sessions.

Data-centric controls are limited but useful for infrastructure access. Database RBAC can restrict access to specific databases, tables, and operations such as SELECT, INSERT, UPDATE, and DELETE. SQL activity is captured in the audit log. For Windows Desktop Access, clipboard and directory sharing can be configured per role. Session recording provides forensic evidence of data access. Crown Jewel tagging allows critical assets to receive enhanced monitoring priority.

Workloads, APIs, service accounts, CI/CD jobs, microservices, and AI agents receive cryptographic identities with assigned roles. Secure MCP enforces Zero Trust access to MCP servers used by AI agents. Agent-to-tool invocations are authenticated, authorized against real-time policy, and logged with full context. The Agentic Identity Framework provides reference designs for long-running agents, LLM applications, delegate workflows, MCP discovery, LLM guardrails, rate limiting, and budgeting.

Teleport holds ISO 27001, SOC 2 Type II, FIPS 140-2, FIPS 197, and PCI DSS certifications. Deployment options include Teleport-managed SaaS, self-hosted environments, and hybrid configurations. Policy-as-code is supported through SDKs, Terraform, Ansible, GitOps workflows, and a Kubernetes Operator. Documentation and support are available in English, with 24/7 support for Severity 1 incidents and professional services available globally.

Teleport is best suited to engineering-forward organizations with complex, distributed infrastructure across clouds, on-premises environments, Kubernetes, databases, and AI workloads. Its certificate-centric, secretless architecture directly addresses credential sprawl, lateral movement risk, and non-human identity governance. The platform is strongest in secure infrastructure connectivity, machine and workload identity, AI agent governance through MCP, and audit depth.

Strengths


- Secretless architecture issues short-lived certificates to all identities at access-request time
- Identifies stolen renewable certificates by detecting mismatches between expected and observed certificate usage counts
- Database RBAC enforces access at the database, table, and operation level, with SQL activity captured in structured audit logs
- Purpose-built AI agent governance for infrastructure access, including identity-aware session recording, approval workflows, and audit trails for agent actions
- Secure authentication and governance for AI agent-to-tool invocations
- AI-powered Session Summaries identify privilege escalations and other security-relevant behaviors

Challenges

- The platform is infrastructure-access-oriented; it is designed to complement rather than replace ZTNA/SSE platforms
- No DLP or content inspection capabilities
- Session risk signals raise alerts but do not currently trigger automated session revocation (planned for the near future)
- Customer install base concentrated in North America; limited EMEA and APAC market presence

ThreatLocker – ThreatLocker Zero Trust Platform

THREATLOCKER



Founded in 2017 and headquartered in Orlando, Florida, ThreatLocker is a privately held cybersecurity company. Its core model is Zero Trust for endpoints, with default-deny application control, allowlisting, privilege elevation control, storage control, network controls, and its “Ringfencing” capability to limit what approved applications can actually touch. ThreatLocker serves organizations across North America, EMEA, and APAC, with a customer base spanning small businesses, mid-market organizations, and large enterprises in sectors such as financial services, manufacturing, healthcare, aerospace and defense, government, and retail.

ThreatLocker’s architecture is built on a centralized cloud-managed control plane delivered through the ThreatLocker Portal and a locally installed agent that caches and enforces policies directly on endpoints. Policy decisions are executed locally in real time, so, during cloud service degradation, cached policies continue to enforce controls without interruption. The platform supports Windows, macOS, Linux across a broad set of distributions, iOS, and Android. Deployment is available as ThreatLocker-hosted SaaS or as a self-managed customer-hosted option, with geographic data center selection available. A mobile application supports the Zero Trust Cloud Access capability.

The foundational principle is deny-by-default, allow-by-exception. During initial deployment, the agent enters Learning Mode, logging application activity without blocking it and generating policies for observed applications before transitioning to Secured Mode. ThreatLocker also maintains a large library of built-in application definitions for common business software, with hashes updated as applications change. This reduces operational disruption from software updates and lowers the administrative burden that historically made allowlisting difficult to manage.

Application Allowlisting is the core enforcement mechanism, controlling which executables, scripts, libraries, and drivers can run. Policies can be defined at organization, computer group, or machine level and scoped to users and groups. Just-in-time access is provided through Elevation Control, which replaces persistent local administrator rights with per-application elevation. Elevation can be combined with Ringfencing to ensure that even elevated applications remain constrained.

Ringfencing is ThreatLocker's most distinctive capability: it controls what trusted applications can do while running. Policies can prevent applications from launching other applications, accessing specific files or file types, modifying registry areas, reaching the internet, or interacting with network shares. For example, Office applications can be blocked from launching tools such as Command Prompt, PowerShell, scripting engines, and other commonly abused Windows components. A monitor paths feature allows Ringfencing to apply to file types regardless of their filesystem location. Cascading Ringfencing is under development to apply a parent application's restrictions recursively to child processes. Ringfencing can also constrain AI-related tools by controlling whether they run at all and what files, processes, or internet destinations they can reach once running.

Network Control provides host-based microsegmentation independently of Windows Firewall. When one device attempts to connect to another, ThreatLocker performs challenge-response validation to confirm that the source device belongs to an authorized object before opening the port. Unauthorized or unmanaged devices receive no response, limiting discovery, scanning, and enumeration. This directly addresses ransomware encryption attacks, where compromised VPN or firewall access is used to reach SMB shares or other internal resources. Multi-group membership for devices is under development to enable more nuanced contextual policy assignment.

ThreatLocker ZTNA provides brokered connectivity for device-to-device and user-to-resource access without exposing VPN ports to the internet. Policies define source objects, destination devices, ports, and protocols, with traffic routed through the ThreatLocker broker. Zero Trust Cloud Access protects SaaS and cloud applications against credential phishing, token replay, and session abuse by routing access through a ThreatLocker-managed secure broker. Services such as Microsoft 365, Salesforce, Snowflake, and Dropbox can be restricted to approved broker IPs. Customers must configure the corresponding conditional access or IP allowlisting controls on the cloud service side.

Storage Control governs data access across local folders, network shares, cloud storage, and USB devices. Policies can restrict storage access by process. File activities are logged with user and process context. USB access is default-deny, with encrypted device

enforcement available. This enables DLP through access restriction, although content inspection or redaction are not provided. Browser session controls, including copy-paste, download, upload, printing, and screenshot restrictions, are supported through policies.

ThreatLocker Detect is the platform's EDR component, providing behavioral telemetry analysis, anomaly detection, and alerting. Policies can be chained to create detections from correlated behaviors rather than isolated signals. The Incident Center provides a workspace for containment and remediation. Automated response actions include endpoint isolation, account lockout, session termination, OS updates, and API or webhook calls to external systems. Cyber Hero MDR provides human-led monitoring and response. Defense Against Configurations identifies security misconfigurations and maps controls to frameworks such as NIST, CMMC, HIPAA, PCI DSS, and SOC 2. Patch Management has also been added to the platform.

Monitoring is centralized through Unified Audit, which records application execution, file access, network communication, elevation events, policy decisions, and user activity across managed devices. An optional syslog ingester can bring firewall, switch, and network equipment logs into the same audit view. Telemetry is exportable through REST APIs, webhooks, and syslog to SIEM and SOAR platforms.

ThreatLocker holds Common Criteria, and SOC 2 Type II for Security certifications and is on the FedRAMP marketplace. REST APIs with authenticated tokens are available for automation, but IaC tooling is not currently supported. Policy-as-code is available through CSV-based configuration. Policy testing can be performed using an application testing environment and Monitor Only mode. The administrative experience is designed around a single portal, single agent, single approval workflow, and single subscription model. Support and documentation are available in English, with multilingual support available on request.

The roadmap includes ongoing platform enhancements such as cascading Ringfencing, multi-group device membership, and expansion of Defense Against Configurations, alongside additional undisclosed capabilities.

ThreatLocker is best suited to organizations with stringent compliance or threat-defense requirements, including regulated enterprises in financial services, healthcare, government, and defense. At the same time, it also serves SMB and mid-market organizations seeking comprehensive endpoint enforcement across application execution, application containment, host-based microsegmentation, cloud access protection, storage governance, web filtering, and MDR from a single agent and management interface. Its strongest differentiators are mature Application Allowlisting, Ringfencing, and host-based network control, especially for ransomware and lateral movement defense.

Strengths


- Ringfencing addresses a lateral movement vector that network-layer Zero Trust cannot prevent

- Agentless enforcement for the majority of segmentation use cases through native OS firewall capabilities
- Unified application allowlisting, storage control, network control, and cloud access protection in a single agent and management interface
- Host-based microsegmentation gates each connection through device authorization before any port opens, blocking lateral discovery without a separate firewall
- Zero Trust Cloud Access protects SaaS applications from credential phishing and token replay by routing access through a managed broker
- Being on the FedRAMP marketplace extends the platform's reach into US federal and regulated-sector deployments

Challenges

- The platform is endpoint-centric by design
- No DLP or data-centric security capabilities beyond access restriction and storage control
- Policy-as-code is limited to CSV-based configuration; IaC tooling (Terraform, GitOps) is not currently supported
- Limited applicability for cloud-native workloads and containers

Unosecur – Unosecur



Founded in 2021 and headquartered in Berlin, Germany, Unosecur is a privately held cybersecurity company focused on identity risk across complex cloud and SaaS environments. The company targets highly regulated European enterprises, where identity sprawl across clouds, SaaS applications, DevOps platforms, and AI systems creates material security exposure. The product evaluated here is the Unosecur platform, a SaaS-based identity security product positioned as a Unified Identity Fabric. Its goal is not to replace identity providers, PAM, IGA, or Zero Trust access platforms, but to connect them and expose identity risk that remains hidden when each system is viewed in isolation.

The Unosecur platform has self-managed options available on request. Deployments are available in EU and India regions, with additional regions available on request. For highly regulated and air-gapped environments, the platform can also be deployed within the customer's own environment, including on-premises and fully air-gapped installations. The platform runs on Kubernetes with horizontal scaling and multiple API service replicas for high availability. Data services run within a private VPC. The platform connects disparate identity systems through native API connectors rather than relying only on standard protocols such as SCIM. This allows Unosecur to collect and correlate identity, permission, and activity data for human users, NHIs, and AI agents across the customer's entire IT landscape.

Instead of relying on periodic snapshots, Unosecur operates as a continuous runtime data stream, ingesting identity and activity signals from connected systems in near real-time. The connector library covers major cloud platforms, identity providers such as Okta, Microsoft Entra ID, and PingFederate, SaaS applications such as Salesforce, GitHub, Jira, Snowflake, and HubSpot, security platforms such as Wiz and Prisma Cloud, PAM tools, HR systems, and ticketing platforms such as ServiceNow and Jira. By using native APIs, Unosecur can expose identity objects and activity that SCIM-based integrations often miss, including locally managed users, shadow administrators, service principals, OAuth applications, and activity logs outside the visibility of standard identity management tools.

The Unified Identity Fabric dashboard aggregates connected identity data into a normalized and correlated view. Human identities are matched across systems using correlation techniques. Supported NHIs include service accounts, API keys, CI/CD tokens, OAuth applications, Kubernetes workloads, and AI agents. The Identity Access Graph visualizes identity-to-permission relationships across systems, showing how access is granted through groups, roles, policies, and inherited entitlements. This helps identify access paths that would not be visible inside any single system.

Unosecur analyzes identity activity from connected systems to determine which permissions, roles, and services each identity actually uses, then generates least-privilege recommendations based on observed behavior. For AWS, this includes generating least-privilege IAM policies that can be deployed directly or exported as Terraform or Ansible playbooks for review through established infrastructure-as-code and change management workflows. For Azure, recommendations cover service principals, managed identities, and custom roles. Customers are encouraged to review proposed changes using dry-run or simulation approaches and route them through ITSM-based approval workflows with four-eyes review. The platform operates read-only by default, with execution permissions required only when customers enable automated remediation.

The Findings engine generates actionable insights from connected identity data. These include users without MFA, unused or unrotated access keys, inactive accounts with valid credentials, service principals with excessive permissions, shadow administrators, etc. Individual signals are composed into higher-confidence risk findings. Built-in compliance mapping supports frameworks such as PCI DSS, NIS2, DORA, SOC 2, and ISO 27001, linking specific findings to remediation steps. Ticketing integrations allow findings to be assigned to owners with full context.

The platform continuously baselines identity behavior and flags deviations such as abnormal authentication times, unusual API activity, data access anomalies, and suspicious privilege escalation. Because detections are based on Unosecur's own activity ingestion, the platform can cover a broad set of applications and identity types. Context enrichment from tools like Wiz and Prisma Cloud allows access paths to be correlated with data sensitivity. HR system integration helps detect high-risk cases such as authentication activity after an employee has left the organization.

Unosecur discovers AI agents and autonomous systems across connected environments, maps their permissions and resource access, and applies the same least-privilege analysis

and behavioral monitoring used for human and non-human identities. An emerging MCP gateway capability is offered for AI agent governance, and AI agent-specific detection scenarios are being added to the threat catalog, including unauthorized resource access and anomalous permission use.

The solution operates post-IdP, observing runtime identity behavior, mapping permissions, identifying toxic access paths, and supporting remediation. SIEM integration for log export and alerting is supported through REST APIs and webhooks. Automated response actions include access revocation, account quarantine, and role downgrades, either from the findings interface or routed through ITSM workflows.

Unosecur holds ISO 27001 certification and SOC 2 Type II attestation across all five trust service criteria. Support is provided in English and German. The roadmap includes expanded AI agent detection, agentic AI workflows, deeper NHI correlation and attribution, more granular internal RBAC, additional identity provider login methods, and platform scalability improvements.

Unosecur is best suited to security teams in regulated European enterprises that need visibility into identity risk across complex multi-cloud and multi-SaaS environments. Its strongest use cases involve locally managed users, service principals, OAuth applications, over-privileged cloud roles, residual access after off-boarding, and emerging AI agent identities. The platform's main strengths are deep native API connectors, activity-based least-privilege enforcement, and a unified correlation layer for fragmented identity records across systems. It is best understood as a specialized identity risk and governance layer that improves the effectiveness of existing Zero Trust and identity infrastructure rather than replacing it.


Strengths

- Native API connectors expose identity objects and activity that standard integrations miss
- Identity Access Graph visualizes user-to-permission and resource-to-user relationships across all connected systems simultaneously, revealing toxic access paths
- Activity-based least-privilege enforcement generates policies from observed behavior rather than static role definitions
- Continuous runtime identity activity stream enables detection of behavioral anomalies cross-referenced with HR systems
- AI agents are discovered, analyzed, and monitored along with human identities within the same unified fabric
- EU-headquartered and EU-hosted, directly addressing European regulatory and data residency requirements
- Native compliance mapping for PCI DSS, NIS2, DORA, SOC 2, and ISO 27001, linking findings directly to remediation steps

Challenges

- The platform is a post-IdP identity risk and governance layer designed to complement, not replace, ZTNA, segmentation, and access enforcement platforms
- Read-only by default; automated remediation (access revocation, account quarantine, role downgrade) is opt-in and routed through ITSM or approval workflows, not executed inline at an enforcement point
- Smaller installed base than established incumbents; market presence concentrated in European enterprises
- Agentless architecture with no inline network- or endpoint-level enforcement; packet-level blocking is delegated to existing enforcement points
- Connector coverage is cloud- and SaaS-centric with native on-premises directory support (Active Directory, LDAP); OT/IoT and legacy systems may still require custom integration
- Internal RBAC is coarse-grained; more granular role separation within the platform is a roadmap item rather than a current capability

Whiteswan – Whiteswan Identity Security Platform



Founded in 2023 and headquartered in San Mateo, California, Whiteswan Identity Security is an early-stage privately held cybersecurity company focused on identity-centric Zero Trust access. The company's vision is based on the premise that identity and access are the most actionable entry points for implementing Zero Trust in practice. Whiteswan primarily serves mid-market organizations in manufacturing, healthcare, pharmaceutical, and aerospace and defense sectors. The solution is still at an early stage of commercial maturity and certification coverage.

The Whiteswan platform is organized as an integrated suite covering four related capability areas: trusted access for on-premises servers and endpoints, cloud identity governance and just-in-time access, SaaS application access with credential vaulting, and an agentic AI gateway for MCP-based AI agent governance, currently in beta. The architectural foundation is a mesh VPN overlay built on Nebula, the open-source peer-to-peer overlay network. Whiteswan extends this with device certificates to enable cryptographically authenticated membership in the mesh. A Lighthouse node, hosted either by the vendor or within the customer's own deployment, coordinates mesh membership and enables NAT traversal.

This allows peer-to-peer connectivity between enrolled devices and resources without opening inbound ports.

Deployment models include vendor-hosted SaaS, self-managed Kubernetes deployment in a customer cloud or on-premises data center, and Docker-based deployment for smaller environments. The platform is multi-tenant with logical tenant isolation, dedicated tenant options, and geographic region selection for tenant data. High availability is implemented through Kubernetes HA for control plane components.

Enforcement relies on a lightweight user-mode agent for Windows, macOS, and Linux endpoints and servers. For network devices and resources that cannot host an agent, a gateway mode allows a Linux agent in the same subnet to act as a conduit, using IP forwarding to provide access through the mesh. Authentication options include SAML, OIDC, Active Directory, LDAP, Entra ID, Okta, token-based authentication, FIDO2, passkeys, Windows Hello, among others. Data in transit is encrypted using TLS 1.3 and the Nebula overlay protocol. Data at rest is encrypted with AES, with customer-managed keys held on-premises.

Access policy supports RBAC, ABAC, contextual conditional access, and risk-based access models. Authorization granularity extends from server and system level through application, service, workload, and API route level down to function and method level for MCP tool invocations. JIT and time-bound access are the core operating models. Administrators publish resources, such as servers, cloud accounts, Kubernetes clusters, or SaaS applications, to user roles with defined access scopes. Users request access through a self-service portal, receive approval, and are granted a time-limited session that is automatically revoked when it expires.

Continuous session re-evaluation is supported, with triggers including identity risk signals, device posture changes, location or network changes, and behavioral anomalies. Approval workflows for sensitive access requests are supported through ServiceNow integration. Step-up authentication applies adaptive MFA not only to application access but also to specific actions such as command-line utility invocation, folder access, privilege elevation, and PowerShell execution. This allows authentication to be enforced at the moment of sensitive action, not just at login.

For on-premises endpoint and server access, Whiteswan discovers accounts and users on enrolled systems and allows administrators to place specific command-line utilities into a denied-by-default state. Examples include utilities like curl, nslookup, scripting tools, and network enumeration commands that are often used by malicious actors. Users can request JIT elevation when they need to invoke these utilities, and all invocations are logged. For SSH access, Whiteswan generates a short-lived key pair per session, with the private key remaining on the server and the public key added to the target's authorized keys file, reducing the burden of manual SSH key management. Remote session recording is provided through Apache Guacamole.

For SaaS application access, Whiteswan vaults credentials in OpenBao (open-source fork of Hashicorp Vault) and uses a browser automation engine to inject credentials into web

applications without exposing them to the user. This is particularly useful for shared or service account credentials in applications that do not support SAML or OIDC federation. The platform automatically fills in login forms for supported applications, and Whiteswan can fine-tune this behavior for applications with unusual or custom login pages.

For cloud identity governance, the platform connects to AWS, Azure, and GCP, discovers human and NHIs, including service accounts, managed identities, and Kubernetes secrets, and ingests activity logs to build behavioral baselines. An IAM analyzer identifies overprivileged roles and unused permissions, generating right-sizing recommendations from observed activity rather than static role definitions. A cloud identity graph visualizes user-to-resource and resource-to-user relationships. JIT cloud access follows the same request, approve, and terminate model used for server access, with time-bound role assignments across public clouds and Kubernetes that are revoked automatically at session end.

Monitoring and analytics cover session logs, policy decisions, identity context, and behavioral anomaly detection. Real-time session visibility is available through the administration console. SIEM integration is supported through syslog and API-based log streaming. Automated responses include session termination, step-up authentication, account and identity blocking, device isolation, and AI-assisted decision support. Policy decisions are traceable to contributing signals, supporting forensic investigation. The ITDR component detects lateral movement attempts, including Golden Ticket attacks, abuse of native PowerShell remoting utilities, and anomalous privilege escalation. Threat signals can dynamically re-trigger access decisions and revoke sessions.

NHI support covers workloads, service accounts, and AI agents. Workload identity federation is supported, and least-privilege enforcement for machine identities follows the same JIT access model used for human identities. The MCP agentic gateway, currently in beta, onboards MCP servers as protected resources and authenticates AI agents through OAuth integrations such as Okta/Auth0. It controls which MCP tools an agent may invoke, and unauthorized tool access can trigger a human-in-the-loop approval request.

Policy-as-code is supported through Rego, with REST APIs, JSON, and Terraform-based integration available. Support is provided through a partner model, with English-language support and documentation. On-site support is available across North America, EMEA, and APAC through regional coverage. The roadmap includes broader identity mesh and risk signal capabilities, automated permission right-sizing, SPIFFE integration for workload identity, workload identity segmentation, and general availability of the MCP agentic gateway as deployments mature.

Whiteswan Identity Security is best suited to mid-market organizations, particularly in manufacturing, pharmaceutical, and regulated industrial sectors, that want to combine privileged access management, JIT cloud access governance, SaaS credential management, and endpoint command-line controls without deploying multiple separate PAM, vault, and VPN products. Organizations evaluating Whiteswan should account for its early commercial maturity, limited certifications, and beta status of some capabilities.


Strengths

- Credential vaulting with browser automation for SaaS applications that do not support identity federation enables Zero Trust access controls for legacy services
- Step-up authentication granularity for specific action invocations (curl, PowerShell, and network enumeration commands)
- Eliminates persistent SSH key sprawl without requiring migration of existing key files
- ITDR component detects infrastructure attack patterns not commonly addressed by network-centric ZTNA platforms
- JIT cloud access for AWS, Azure, GCP, and Kubernetes follows the same request/approve/terminate model as server access
- Nebula overlay mesh with device certificates allows direct peer-to-peer connectivity for all devices, even across NAT boundaries

Challenges

- Early-stage commercial maturity with limited security certification coverage
- The platform addresses infrastructure, server, cloud, and SaaS access but not web or SaaS traffic inspection
- Support is partner-delivered; direct vendor support availability may be constrained for organizations requiring enterprise SLAs
- Limited market presence; primarily established in manufacturing, healthcare, pharmaceutical, and industrial sectors in North America

Zero Networks – Zero Networks Platform



Founded in 2019 and headquartered in Tel Aviv, Israel, Zero Networks is a privately held cybersecurity company focused on automated microsegmentation, identity segmentation, and secure remote access. The company's vision is rooted in a pragmatic observation: microsegmentation projects usually fail not because the objective is wrong, but because the manual effort required to build and maintain rules is unsustainable. Zero Networks addresses this by using deterministic automation to study the environment, generate policies, and enforce them through native operating system controls.

The Zero Networks platform integrates four main capabilities under a unified console and shared automation engine: NetSeg for network microsegmentation, IDSeg for identity

segmentation of service and administrator accounts, RPC Firewall for granular RPC protocol enforcement, and ZTNA Connect for secure remote access. The control plane is delivered as a multi-tenant SaaS service on Google Cloud Platform, with customer isolation and geographic data residency options. Enforcement is handled by a hardened Segment Server, a Windows virtual machine deployed in the customer environment. It connects to managed assets using native protocols: WinRM and PowerShell for Windows hosts, SSH with iptables or nftables for Linux, and the native firewall on macOS.

For most use cases, no endpoint agent is required: enforcement is applied through built-in OS firewall capabilities, with the Segment Server pushing rules remotely via WinRM for Windows, SSH for Linux, and native management protocols for macOS. An optional agent is available for roaming devices and air-gapped scenarios. Enforcement continues locally if the management plane is temporarily unavailable, since the underlying OS firewall rules remain active. The Segment Server can be scaled horizontally by adding additional instances.

During the monitoring stage, Zero Networks observes network activity from enrolled assets and forwards metadata to the control plane. During learning, the engine analyzes traffic patterns to identify data flows, classify assets using application fingerprints, distinguish client-to-server from server-to-server communication, and generate least-privilege firewall recommendations. During segmentation, generated rules are pushed to assets through native OS controls. Traffic is blocked unless explicitly permitted by generated or administrator-defined rules. Privileged ports such as RDP, WinRM, SSH, and SMB are treated differently: rather than opening them broadly during learning, the platform generates JIT MFA policies for them (described in detail below). Administrators can require approval before specific rules are enforced, while allowing lower-risk rules to be applied automatically.

Simulate Segmentation tests observed activity against the current rule set and reports what would have been allowed, blocked, or MFA-challenged without changing the live environment. This supports validation for individual assets or entire groups before enforcement. In addition, the platform offers enhanced visualization and policy simulation, including asset discovery, communication flow mapping, and natural-language investigation of network behavior, exposure, and access paths.

Zero Networks' most distinctive technical innovation is JIT network-layer MFA. Instead of relying only on application-layer MFA, the platform keeps privileged ports closed at the OS level. When a user attempts to connect to a privileged port, the connection is initially blocked. Zero Networks detects the attempt, evaluates it against policy, identifies the user through the connected identity provider, and triggers MFA. Only after successful authentication, a temporary firewall rule is applied, opening the port just for that user's source machine and only for the approved session window. This approach reduces the exploitability of exposed privileged services even where valid credentials have been compromised. The mechanism applies to Windows, Linux, macOS, and Kubernetes administrative access. For OT devices that lack manageable host firewalls, Zero Networks integrates with Cisco switches to enforce Layer 2 ACLs. An optional lightweight agent serves roaming and air-gapped endpoints that the Segment Server cannot reach via remote management protocols.

IDSeg addresses service accounts and administrative accounts, two common lateral movement paths in Windows environments. The automation engine connects to Active Directory, identifies service accounts, observes their logon behavior and target systems, and generates least-privilege local logon policies. These restrict each service account to the logon types and systems it actually needs. For administrator accounts, policies can require MFA before privileged server access and prevent domain admins from interactively logging into client systems where credential material might be cached.

RPC Firewall addresses a specific Windows domain controller problem. Traditional microsegmentation often struggles with RPC because domain controllers require many dynamic ports for normal operation. Zero Networks inspects the RPC interface and operation from connection payloads and enforces policies at that level. This allows legitimate domain activity while reducing exposure to remote process execution and RPC-based privilege escalation.

ZTNA Connect provides secure remote access without exposing VPN concentrators. A Secure Remote Access server is deployed on-premises or in the cloud with inbound ports closed by default. Users authenticate to the Zero Networks cloud first, after which the access server establishes a direct WireGuard tunnel between the user client and protected server. Traffic does not pass through the Zero Networks cloud. Device posture signals from compatible SAML identity providers can be included in the authentication flow. A browser-based web portal allows third-party contractors to access specific applications after MFA without requiring VPN software.

Monitoring and analytics functions review session logs, network flows, policy decisions, and identity context. Pre-built views highlight RDP users, internet-facing systems, SSH connections, and other common exposure categories. Natural language querying helps administrators investigate access paths. Telemetry can be exported through syslog, log streaming, and REST APIs to SIEM. ServiceNow integration supports ITSM workflows. Automated responses include session termination, step-up MFA, identity blocking, device quarantine, and real-time revocation of JIT rules. A ransomware auto-blocker can isolate assets showing lateral movement or ransomware-like behavior.

Zero Networks does not perform TLS termination, DLP, or broad protocol inspection beyond RPC. Its value lies in reducing reachable attack surface through default-deny segmentation, closing privileged ports until MFA is completed, limiting service account behavior, and blocking unauthorized east-west access. URL-aware DNS hooking extends firewall enforcement to domain-based policies, while process-level firewall rules can allow or block applications. Kubernetes support covers OpenShift via CNI integration as well as managed Kubernetes environments.

The company holds ISO 27001, SOC 2 Type II, and FIPS 140-2 certifications. A full REST API is available for programmatic access, with Terraform, Python, and PowerShell SDKs for policy-as-code workflows. Support and documentation are provided in English, with 24/7 support options and global on-site support available.

Zero Networks is best suited to organizations that need to reduce east-west lateral movement, eliminate standing privileged network access, and control service account behavior, especially where earlier microsegmentation projects failed due to complexity. Its primary differentiator is automated policy generation combined with just-in-time network-layer MFA for privileged ports.

Strengths

- Fully agentless microsegmentation for the majority of use cases; no endpoint agent is required on enrolled servers
- MFA is gated at the OS firewall layer, with privileged ports kept closed until authentication succeeds
- Addresses service and administrator account lateral movement: a targeted control for the most common ransomware propagation paths
- Granular enforcement at the RPC interface for domain controllers, blocking remote execution and privilege escalation while permitting legitimate domain activity
- Automated policy generation from observed network traffic eliminates the primary failure mode of manual microsegmentation projects
- Natural language network behavior investigation lowers the operational barrier for security teams without network expertise

Challenges

- The platform does not provide inline TLS termination, DLP, or full SWG/CASB inspection of SaaS and internet traffic
- Device posture context relies on SAML identity provider signals; direct integrations with EDR, XDR, and MDM platforms are not natively supported
- No native user behavioral analytics or anomaly detection beyond the ransomware auto-blocker; UEBA-style monitoring depends on external SIEM and EDR integrations
- Primarily a Windows-Active-Directory-centric platform; Linux, macOS, and cloud-native environments receive progressively less coverage

Vendors to Watch

Besides the vendors covered in detail in this document, we observe some other companies in the market that readers should be aware of. These vendors did not participate in the rating for various reasons but, nevertheless, offer a significant contribution to the market space.

1Password

Founded in 2005 and headquartered in Toronto, Canada, 1Password expanded from consumer and enterprise password management into a broader access platform. 1Password Unified Access is the company's Zero Trust offering, combining Enterprise Password Manager with SaaS Manager (from the 2025 Trelica acquisition), Device Trust (from the 2024 Kolide acquisition), Credential Broker for machine workloads, and Privileged Access through the recent Apono acquisition.

Why worth watching: 1Password is pursuing convergence across access management, identity governance, privileged access, and non-human identity from a vault-centric foundation, in contrast with the static-infrastructure focus of legacy PAM providers. The Apono acquisition adds zero standing privilege and runtime-governed just-in-time access for cloud, databases, and infrastructure, with extension to AI agents in beta, while the Credential Broker delivers scoped, short-lived secrets for CI/CD pipelines, Kubernetes, and cloud workloads. Integration of these recent acquisitions remains in progress, so the unified platform is still consolidating. The combination is particularly relevant for organizations pursuing identity-first Zero Trust architectures that span humans, machines, and AI agents.

Absolute Security

Founded in 1993 and headquartered in Seattle, Washington, Absolute Security is a privately held company owned by Crosspoint Capital Partners, rebranded from Absolute Software following its acquisition in 2022. The company's Zero Trust portfolio builds on the acquisition of NetMotion in 2021, which became the foundation of the Absolute Secure Access product line. The portfolio includes ZTNA, VPN, Secure Web Gateway (SWG), CASB, DLP, Remote Browser Isolation (RBI), and Digital Experience Management (DEM), delivered as integrated capabilities within the Absolute Enterprise SSE platform. A distinctive aspect of the company's approach is the persistence technology embedded in endpoint firmware through OEM partnerships, enabling its self-healing client functionality.

Why worth watching: Absolute Security differentiates itself through its focus on endpoint-resident enforcement and persistence. Its self-healing client technology can automatically restore or reinstall critical security components if they are tampered with or removed, helping maintain policy enforcement under adverse conditions. Absolute places significant enforcement logic on the endpoint itself, which can reduce latency and maintain visibility across both tunneled and non-tunneled traffic. This architecture is particularly relevant for mobile and distributed workforces, as well as organizations operating in regulated sectors

such as healthcare, government, and defense, where device resilience and continuous policy enforcement are operational priorities.

Akamai Technologies

Founded in 1998 and headquartered in Cambridge, Massachusetts, Akamai Technologies is a global provider of content delivery, cloud, and security services. Its Zero Trust-related portfolio includes Enterprise Application Access (EAA) for identity-aware application access, Enterprise Threat Protector (ETP) for SWG functionality, and Akamai Guardicore Segmentation for microsegmentation across hybrid and multi-cloud environments following the acquisition of Guardicore in 2021. These capabilities are integrated within the Akamai Connected Cloud platform alongside WAAP, API security, and bot management offerings.

Why worth watching: Akamai's position in the Zero Trust market is shaped by the combination of its global edge infrastructure and its Guardicore microsegmentation technology. While many ZTNA offerings focus primarily on securing user-to-application access, Akamai extends policy enforcement across multiple layers, including edge access brokering, east-west segmentation, and API and application protection. This combination enables organizations to address both initial access control and lateral movement risks within complex hybrid environments. The integration of networking, application delivery, and security functions under a common platform may be particularly relevant for large enterprises with distributed infrastructures and extensive internet-facing application footprints.

Aryaka

Founded in 2009 and headquartered in Santa Clara, California, Aryaka is a privately held company focused on managed SASE services. Its Universal ZTNA offering forms part of the Aryaka Unified SASE platform, integrating ZTNA with SWG, FWaaS, CASB, Next-Generation DLP, SD-WAN, and observability capabilities. The ZTNA functionality incorporates technology originating from Cloudbrink and operates across Aryaka's global private backbone infrastructure, which includes more than 100 PoPs worldwide. The platform also includes AI-focused governance capabilities intended to manage and monitor traffic related to generative AI applications.

Why worth watching: Aryaka emphasizes the operational simplicity of a fully managed service model combined with a private WAN backbone rather than relying exclusively on the public internet. By integrating networking and security services into a single delivery architecture, the company aims to address both connectivity performance and policy enforcement simultaneously. Its single-pass inspection model applies security controls across ZTNA, SWG, DLP, and firewall services within a unified traffic flow, reducing architectural complexity and limiting repeated inspection overhead. Aryaka's approach may appeal particularly to organizations with globally distributed users and a limited appetite for managing multiple networking and security platforms internally.

Barracuda Networks

Founded in 2003 and headquartered in Campbell, California, Barracuda Networks is a privately held cybersecurity company with a longstanding presence in email security, network security, and data protection markets. Its SecureEdge platform represents the company's move toward a converged SASE architecture, integrating ZTNA, Secure Web Gateway (SWG), Firewall-as-a-Service (FWaaS), SD-WAN, and centralized security management into a cloud-delivered platform. It combines lightweight edge appliances, endpoint agents, and cloud-native security services to provide identity-aware access to private applications and internet resources across distributed environments. The platform is closely integrated with Barracuda's broader portfolio, including email protection, XDR, and managed security services.

Why worth watching: Barracuda's relevance in the Zero Trust market stems primarily from its focus on operational simplicity and mid-market accessibility. Rather than targeting highly customized large-enterprise Zero Trust architectures, Barracuda emphasizes unified deployment, centralized management, and ease of adoption for organizations with limited security and networking resources. The SecureEdge platform combines networking and security controls within a single cloud-managed framework, reducing the operational complexity often associated with multi-vendor SASE implementations. Barracuda also maintains a strong channel and managed service provider ecosystem, making the platform particularly relevant for organizations consuming Zero Trust capabilities through managed services rather than direct in-house operation.

Cato Networks

Founded in 2015 and headquartered in Tel Aviv, Israel, Cato Networks is a privately held company closely associated with the emergence of the SASE market. Its Cato SASE Cloud platform combines SD-WAN, ZTNA, SWG, CASB, FWaaS, and MDR capabilities into a cloud-native architecture delivered through a globally distributed private backbone. ZTNA functionality underpins both remote user access and site-to-site connectivity within the platform. Cato has been recognized in multiple KuppingerCole Analysts Leadership Compass reports for ZTNA and SASE-related capabilities.

Why worth watching: Cato's architecture was designed as a converged platform from the outset, rather than integrating separately acquired networking and security products. This unified approach enables centralized policy definition and consistent enforcement across multiple traffic types and deployment models. The platform's single-pass inspection model and shared data lake combine network, security, and user telemetry into a common analytics environment, supporting operational visibility across distributed infrastructures. Cato's integrated operational model is particularly relevant for organizations seeking to consolidate networking and security operations while reducing the complexity associated with multi-vendor SASE environments.

Citrix Systems (Cloud Software Group)

Founded in 1989, Citrix Systems became part of Cloud Software Group following its acquisition in 2022. Citrix Secure Private Access provides identity-aware access to private and web applications without exposing the broader network. The product integrates with Citrix's wider portfolio, including Citrix DaaS, NetScaler, and Citrix Analytics. Citrix retains a substantial installed base in enterprise and regulated environments, particularly among organizations using virtual desktop and application delivery technologies.

Why worth watching: Citrix Secure Private Access is closely aligned with the company's established application virtualization and remote access ecosystem. Organizations already using Citrix for desktop and application delivery can extend Zero Trust principles to these environments without introducing separate access platforms. The solution supports adaptive access policies based on identity, device posture, and contextual risk signals, while also supporting clientless access scenarios for unmanaged devices. Citrix's long-standing presence in secure application delivery continues to make it relevant for enterprises seeking to modernize legacy remote access architectures while preserving operational continuity.

Cloudflare

Founded in 2009 and headquartered in San Francisco, California, Cloudflare operates one of the world's largest distributed internet infrastructures. Cloudflare One is the company's SASE and Zero Trust platform, with Cloudflare Access serving as its ZTNA component. The broader platform includes SWG, CASB, DLP, Browser Isolation, Magic WAN, and AI Gateway capabilities. Cloudflare has also been recognized by KuppingerCole Analysts across several adjacent security categories, including API security and email security.

Why worth watching: Cloudflare's Zero Trust services are delivered directly through its globally distributed Anycast network, enabling enforcement close to the user without requiring dedicated regional gateway infrastructure. This architecture allows the company to combine access control, DNS services, DDoS protection, CDN delivery, and application security within a common operational platform. The addition of AI Gateway extends policy enforcement to outbound LLM and AI API traffic, reflecting growing enterprise interest in governing employee interaction with public AI services. Cloudflare's flexible consumption model and infrastructure-light deployment approach may be particularly relevant for cloud-first organizations and smaller teams seeking rapid implementation.

Ergon Informatik (Airlock)

Founded in 1984 and headquartered in Zurich, Switzerland, Ergon Informatik is a privately held, employee-owned software company with a strong presence in the DACH market. Its Airlock Secure Access Hub integrates application-layer ZTNA, WAAP, API security, and CIAM capabilities into a modular access control platform. Key components include Airlock Gateway, Airlock Microgateway, Airlock IAM, and Airlock Anomaly Shield. Ergon has

appeared in numerous KuppingerCole Analysts reports across API security, passwordless authentication, CIAM, and policy-based access management.

Why worth watching: Airlock takes a unified approach to identity, application access, and API protection rather than treating these as separate technology domains. The platform supports both cloud-native and traditional on-premises deployment models, which may be important for organizations with strict sovereignty, compliance, or infrastructure requirements. Its modular architecture is particularly relevant for regulated industries that require flexible deployment choices and strong integration between IAM and application protection capabilities. Ergon's longstanding presence in sectors such as banking, insurance, and public administration has also given the company significant experience in addressing regional compliance and operational requirements.

Forum Systems

Founded in 2001 and headquartered in Boston, Massachusetts, Forum Systems is a privately held company focused on API security and high-assurance access enforcement. Forum Sentry is the company's primary platform, combining API security, identity-aware access control, token management, and schema validation capabilities. The platform is deployed in environments including government, defense, critical infrastructure, and financial services, where high assurance and strict operational controls are required.

Why worth watching: Forum Sentry is notable for its self-contained architecture, designed without reliance on third-party libraries or open-source software dependencies. This design approach addresses supply chain security concerns that are increasingly relevant in sensitive and regulated environments, including classified and air-gapped deployments. The platform also extends into AI governance use cases, including validation and policy enforcement for AI-driven workflows and machine identities. These capabilities position Forum Systems as particularly relevant for organizations with stringent sovereignty, assurance, and operational integrity requirements.

Google

Founded in 1998 and headquartered in Mountain View, California, Google is a wholly owned subsidiary of Alphabet Inc. Its Zero Trust portfolio centers on BeyondCorp Enterprise, a cloud-delivered ZTNA and context-aware access platform based on Google's internal BeyondCorp architecture, which preceded the broader industry adoption of Zero Trust by nearly a decade. The portfolio also extends into workload and service identity through Google Cloud IAM, Workload Identity Federation, and Zanzibar-derived authorization concepts embedded in Google Cloud's policy enforcement.

Why worth watching: Google is notable because BeyondCorp Enterprise is not a retrofitted VPN replacement but the productization of an access model Google originally developed to secure its own global environment over untrusted networks. In the broader Zero Trust Platforms context, Google's relevance also comes from its native policy and identity capabilities for cloud workloads. Workload Identity Federation reduces dependence on static

credentials for service-to-service access, while Google's Zanzibar lineage points toward fine-grained, relationship-based authorization. For organizations already standardized on Google Cloud, these native controls form an important enforcement layer that any platform strategy must build on, integrate with, or at least account for.

Jamf

Founded in 2002 and headquartered in Minneapolis, Minnesota, Jamf focuses on Apple device management and security. Its Zero Trust offering, Jamf Private Access, forms part of the Jamf Trust platform alongside Jamf Pro, Jamf Protect, and Jamf Connect. The platform delivers identity- and device-aware access control while leveraging telemetry from Jamf's endpoint security and device management tools.

Why worth watching: Jamf occupies a specialized position in the Zero Trust market through its focus on Apple-centric environments. The platform integrates endpoint management, identity, posture assessment, and secure access within a workflow designed specifically around macOS and iOS ecosystems. By using native Apple security frameworks and Jamf telemetry, the platform can provide device context and posture signals tailored to Apple deployments. This approach is particularly relevant for organizations operating predominantly on Apple hardware, including sectors such as media, technology, higher education, and life sciences.

Palo Alto Networks

Founded in 2005 and headquartered in Santa Clara, California, Palo Alto Networks is a major cybersecurity vendor with a broad enterprise security portfolio. Prisma Access delivers cloud-based ZTNA and SASE capabilities, including SWG, CASB, FWaaS, and SD-WAN services through a globally distributed infrastructure. The platform integrates with the broader Palo Alto Networks ecosystem, including Cortex XDR, Cortex XSOAR, and AI Runtime Security offerings.

Why worth watching: Prisma Access reflects Palo Alto Networks' broader platform strategy, where telemetry and policy enforcement are shared across network, endpoint, cloud, and automation layers. Integration with Cortex XDR and XSOAR enables access decisions to incorporate endpoint telemetry and automated response workflows, extending ZTNA beyond isolated access control. The addition of AI workload protection capabilities also reflects growing enterprise requirements around AI agent governance and machine identity security. Prisma Access is particularly relevant for organizations already invested in the broader Palo Alto Networks ecosystem and seeking tighter integration between Zero Trust access and security operations.

Sangfor Technologies

Founded in 2000 and headquartered in Shenzhen, China, Sangfor Technologies is a cybersecurity and IT infrastructure vendor with operations across Asia-Pacific, the Middle

East, Europe, and Latin America. Its Zero Trust Guard platform is delivered as part of the Sangfor Athena SASE portfolio, integrating ZTNA, SWG, FWaaS, EDR, SD-WAN, DLP, and data protection capabilities.

Why worth watching: Sangfor represents an important regional perspective within the broader Zero Trust market, particularly across Asia-Pacific and the Middle East. The company combines adaptive access control, endpoint telemetry, malware detection, and agentless access support within a unified platform architecture. Its growing international presence reflects increasing demand for regional alternatives in the SASE and ZTNA space. Sangfor may be particularly relevant for organizations with substantial operations in APAC regions where the company has established channel ecosystems, localized support, and market familiarity.

Skyhigh Security

Spun out from McAfee Enterprise in 2022 by Symphony Technology Group and headquartered in San Jose, California, Skyhigh Security traces its lineage to Skyhigh Networks (founded 2011) and the cloud security business of McAfee Enterprise. The company operates a Security Service Edge platform built organically on a single Skyhigh Cloud Platform, combining Secure Web Gateway, CASB, Universal ZTNA for cloud and on-prem private access, DSPM, and an Enterprise Browser layer delivered via patented JavaScript injection rather than a separate browser product.

Why worth watching: Skyhigh's relevance to Zero Trust comes from an architecture led by data security in which DLP, classification, and DSPM operate as a common layer across web, SaaS, private apps, email, IaaS, and endpoints rather than as separate controls in each channel. Universal ZTNA extends the same policy and classification engine to private application access in cloud and on-prem environments, with shared data fingerprinting across CASB, SWG, and ZTNA. Recent work in AI security applies the same backbone to prompt controls and shadow AI discovery for sanctioned platforms such as Microsoft Copilot and ChatGPT Enterprise. The combination is most relevant for organizations that view Zero Trust through a data-centric rather than a network-centric lens.

Tailscale

Founded in 2019 and headquartered in Toronto, Canada, Tailscale is a privately held company delivering a WireGuard-based mesh networking platform designed to simplify secure connectivity. Unlike traditional gateway-centric ZTNA products, Tailscale establishes encrypted peer-to-peer tunnels directly between devices and services, with centralized policy coordination and identity-aware access control. The platform supports a broad range of environments, including user devices, servers, containers, Kubernetes workloads, and IoT systems.

Why worth watching: Tailscale's architecture avoids many of the centralized gateway bottlenecks associated with traditional ZTNA designs by relying on direct device-to-device connectivity coordinated through a central control plane. This model reduces dependency on

dedicated infrastructure and minimizes traffic hairpinning, while maintaining identity-based policy enforcement. The platform's developer-oriented tooling, infrastructure-as-code integrations, and Kubernetes support have contributed to adoption among engineering and cloud-native teams. At the same time, expanding enterprise capabilities such as SCIM provisioning, centralized logging, and compliance features are broadening its relevance beyond developer-centric deployments.

Twingate

Founded in 2019 and headquartered in San Francisco, California, Twingate is a privately held company focused on software-based, resource-centric ZTNA. Its platform uses a split architecture in which a cloud-hosted control plane coordinates authentication and policy decisions while traffic flows directly between users and lightweight connectors deployed in customer environments. The solution integrates with major identity providers and endpoint posture systems.

Why worth watching: Twingate emphasizes rapid deployment and operational simplicity. Its lightweight connector architecture avoids the need for inbound firewall changes and can be deployed with minimal infrastructure modification. The platform's resource-centric access model enforces least-privilege access at the application and service level rather than exposing broader network segments. This design may be particularly attractive for mid-market organizations, SaaS providers, and technology companies seeking to replace legacy VPN infrastructures without extensive network redesign projects.

Versa Networks

Founded in 2012 and headquartered in Santa Clara, California, Versa Networks is a privately held company delivering integrated SASE and SD-WAN platforms. Versa ZTNA forms part of the broader Versa SASE platform, which combines SD-WAN, SSE, SWG, CASB, NGFW, analytics, and Zero Trust access within a common operating system architecture.

Why worth watching: Versa differentiates itself through its use of a unified operating system for both networking and security functions, rather than integrating multiple acquired products. This architecture enables shared policy enforcement and telemetry across SD-WAN and Zero Trust services, which can simplify operational management for distributed enterprise environments. Versa also supports multiple deployment models, including cloud-managed, on-premises, and MSP-delivered approaches, making it relevant for organizations with sovereignty requirements or service-provider-led delivery models. Its strong presence in the carrier and managed services ecosystem further extends its reach into mid-enterprise deployments.

Zscaler

Founded in 2007 and headquartered in San Jose, California, Zscaler is a major provider of cloud-delivered security services. Zscaler Private Access (ZPA) delivers proxy-based, identity-aware access to applications without exposing the underlying network. ZPA operates as part of the broader Zscaler Zero Trust Exchange platform, which also includes internet access, digital experience monitoring, deception technologies, and workload communication controls.

Why worth watching: Zscaler was among the earliest vendors to operationalize large-scale proxy-based ZTNA, replacing network-centric remote access with application-specific brokering. Its architecture removes direct network exposure by connecting users only to authorized applications through broker services, without requiring inbound firewall exposure. The company has continued extending these principles into workload communications and non-human identity scenarios, reflecting broader Zero Trust adoption beyond user access. Integration with deception technologies and identity threat detection capabilities also expands the platform's role from access enforcement toward broader threat visibility and response.

Related Research

[Leadership Compass: Zero Trust Network Access](#)

[Leadership Compass: Policy-Based Access Management](#)

[Buyer's Compass: Zero Trust Network Access \(ZTNA\)](#)

[Buyer's Compass: Secure Access Service Edge \(SASE\)](#)

[Buyer's Compass: Security Service Edge](#)

[Buyer's Compass: Policy Based Access Management](#)

[Advisory Note: Maturity Level for Zero Trust: A Comprehensive Analysis](#)

[Advisory Note: From Machine Identity to Agentic AI – Charting the NHI Continuum](#)

[Advisory Note: Machine Identities](#)

[Leadership Brief: Zero Trust Platforms](#)

[Whitepaper: From AI Agents to Trusted Digital Workers](#)

[Whitepaper: From Perimeter to Persona: Why Data Security Now Starts with Identity](#)

Copyright

© 2026 KuppingerCole Analysts AG. All rights reserved. Reproducing or distributing this publication in any form is prohibited without prior written permission. The conclusions, recommendations, and predictions in this document reflect KuppingerCole Analysts' initial views. As we gather more information and conduct more profound analysis, the positions presented here may undergo refinement or significant changes. KuppingerCole Analysts disclaims all warranties regarding the completeness, accuracy, and adequacy of this information. Although KuppingerCole Analysts' research documents may discuss legal issues related to information security and technology, we do not provide legal services or advice, and our publications should not be used as such. KuppingerCole Analysts assumes no liability for errors or inadequacies in the information contained in this document. Any expressed opinion may change without notice. All product and company names are trademarks™ or registered® trademarks of their respective holders. Their use does not imply any affiliation with or endorsement by them.

KuppingerCole Analysts supports IT professionals with exceptional expertise to define IT strategies and make relevant decisions. As a leading analyst firm, KuppingerCole Analysts offers firsthand, vendor-neutral information. Our services enable you to make decisions crucial to your business with confidence and security.

Founded in 2004, KuppingerCole Analysts is a global, independent analyst organization headquartered in Europe. We specialize in providing vendor-neutral advice, expertise, thought leadership, and practical relevance in Cybersecurity, Digital Identity & IAM (Identity and Access Management), Cloud Risk and Security, and Artificial Intelligence, as well as technologies enabling Digital Transformation. We assist companies, corporate users, integrators, and software manufacturers to address both tactical and strategic challenges by making better decisions for their business success. Balancing immediate implementation with long-term viability is central to our philosophy.

For further information, please contact clients@kuppingercole.com.