The CNAPP Buyer's Guide



paloalto NETWORKS | CORTEX® CLOUD

INTRODUCTION

Cloud security has reached an inflection point. Organizations face exponentially growing attack surfaces across multicloud environments while adversaries leverage AI to accelerate and sophisticate their campaigns. Traditional approaches that worked for on-premises infrastructures fail in dynamic cloud environments where workloads scale elastically, containers live for seconds, and identity sprawl can create massive risk exposure.

The cloud-native application protection platform (CNAPP) market promised to solve cloud security fragmentation by unifying capabilities into single platforms. Most CNAPPs, though, recreate the same problems—thousands of disconnected findings leaving gaping holes between detection and prevention, and between investigation and response.

Modern cloud security demands more than consolidated dashboards. Organizations need platforms that transform raw telemetry into actionable intelligence. Security can no longer be reactive when infrastructure changes by the second and attacks unfold in minutes.

What Makes a Platform Unified

A unified security platform isn't just multiple tools sharing a user interface. Unification requires a data architecture designed, from day one, for correlation, using normalized telemetry that enables cross-domain intelligence rather than fractured visibility. When vulnerability data, identity context, runtime behavior, and infrastructure configuration operate on the same object model, security teams can finally see complete attack narratives.

The questions in this guide—all real-world questions that security and technology leaders have asked Palo Alto Networks—look beyond surface-level capabilities. They examine the architectural foundations that determine long-term platform value: Can your CNAPP scale without introducing friction? Does it reduce operational overhead or merely centralize complexity? Can it prevent recurring issues rather than just repeatedly detecting them?

How to Use This Guide

Each section addresses critical evaluation areas that distinguish effective CNAPPs from repackaged point solutions. Questions are selected to reveal architectural competency, operational efficiency, and measurable security outcomes.

Evaluate our responses against your organization's cloud maturity, security team structure, and operational constraints. The most sophisticated platform delivers no value if it doesn't align with your deployment realities and staffing capabilities.

Focus on the answers that demonstrate concrete differentiation rather than generic capabilities. Every vendor claims unified visibility and integrated workflows. Probe for specific architectural decisions that enable those outcomes and the quantifiable improvements they deliver.

Cloud security effectiveness ultimately depends on the platforms that match the speed and complexity of modern infrastructure while reducing the operational burden on security teams. Use this guide to identify solutions that transform cloud security from a reactive scramble into a proactive, intelligence-driven discipline.

Table of Contents

I. PLATFORM ARCHITECTURE 5

Can your CNAPP manage the complexity of cloud and security telemetry and turn it into meaningful risk insights? 5

How does your platform use AI to improve security outcomes and reduce manual effort? 8

How does your CNAPP support secure, role-based access and customizable dashboards without fragmenting visibility? 9

What architectural choices simplify onboarding and accelerate time to value? 10

If we're not ready to deploy everything at once, can your platform support staged adoption? 10

II. APPLICATION SECURITY 12

Can you consolidate findings from existing AppSec tools into a single risk, policy, and automation engine? 12

Can your CNAPP show the full chain from repository to pipeline, to image to runtime? 14

What policy-as-code options exist for targeted blocking or reporting in PRs and builds? 15

How does your CNAPP reduce friction between developers and security teams? 17

III. CLOUD POSTURE AND IDENTITY SECURITY 18

How does your CNAPP help teams prioritize and remediate risks across cloud infrastructure? 18

Does your CNAPP preserve context and evidence from short-lived cloud assets to support forensic investigations? 22

How does your platform support compliance, audit readiness, and data residency across global regions? 23

IV. RUNTIME PROTECTION AND RESPONSE 25

What runtime protection options are available, and how do they cover modern compute types? 25

How does your platform detect evasive or unknown threats in real time? 27

MITRE ATT&CK Validation 29

V. SOC INTEGRATION AND OPERATIONAL SCALE 30

How does your CNAPP turn code-to-cloud risk and SOC operations into one continuous motion across every security team? 30

What cost, complexity, or staffing efficiencies can we anticipate if we go with your platform? 32

1 PLATFORM ARCHITECTURE

Can your CNAPP manage the complexity of cloud and security telemetry and turn it into meaningful risk insights?

Security programs fail when tooling assumes the cloud is uniform. In reality, no two cloud environments are the same. One account runs VMs and containers. Another leans on serverless and managed datastores. And, most organizations operate in hybrid states, to say nothing of those living in private cloud deployments layered with custom orchestration. Add in region-specific regulations and third-party tools, and any security platform that lacks architectural flexibility becomes a liability.

Managing that complexity, and the volume of telemetry it produces, requires more than a distributed agent model or surface-level integrations. It demands an architecture that decouples security control from cloud ownership, abstracts telemetry into normalized objects, and centralizes risk logic regardless of where the workload executes. Without it—without a unified data layer engineered from day one—context is nonexistent, and prevention is lost in the inability to prioritize risk.

Cortex® Cloud™ was architected to support deep data sharing from the start.

Unified Data Architecture

Cortex Cloud supports agentless ingestion of configuration, identity, and workload metadata directly via APIs across major cloud providers. For runtime telemetry and enforcement, it deploys lightweight agents and Kubernetes connectors into private or hybrid environments, including on-premises data centers and self-managed clusters.

All events, findings, and assets enter a single data lake under a common schema, stitched together with consistent terminology and asset IDs. Even if multiple scanners analyze the same asset, it appears as one unified object in the platform. Every part of Cortex Cloud can pull from and contribute to this unified data lake. Some data is stored in traditional relational databases for complex search and investigation. Other data is stored in a graph database that maps relationships between assets and findings.

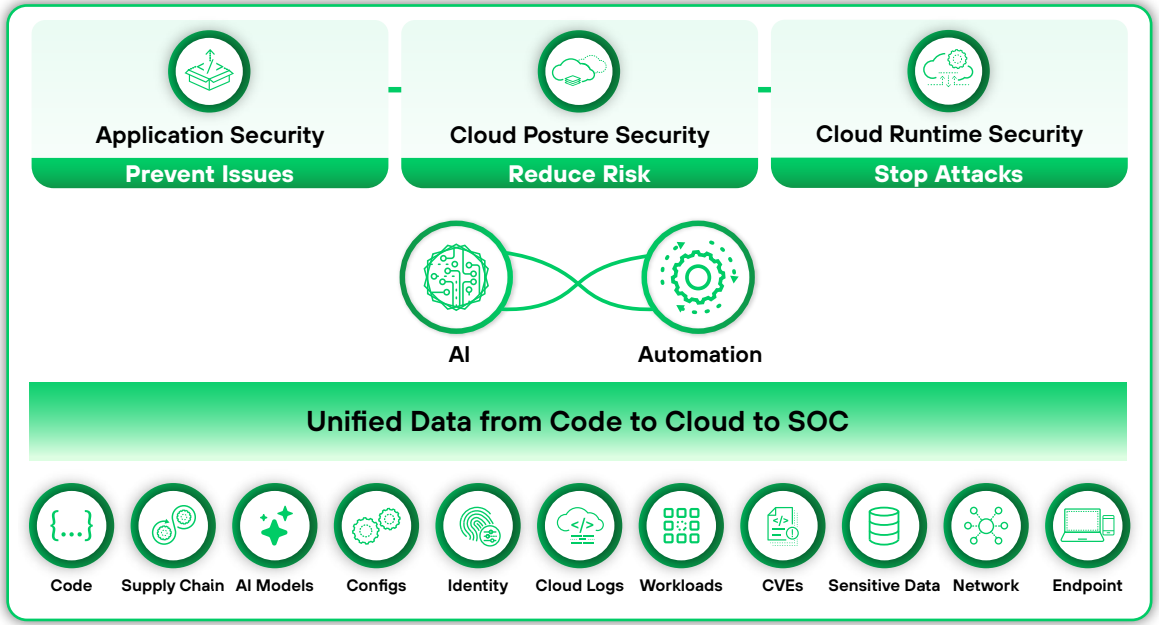


Figure 1. The unified data layer turns telemetry into prioritized risk decisions

Context-Driven Intelligence

Cloud security hinges on context. What matters isn't just whether a workload has a public IP, but whether it's accessible, whether a package is in use, and whether a firewall protects it at the edge. Because Cortex Cloud feeds all telemetry into a single data lake, it can enforce consistent controls and create unified timelines across disparate infrastructure.

An attack detected in a public cloud workload, for example, can be linked to permissions managed through a private cloud IAM bridge or a vulnerability introduced in an on-premises registry. The platform doesn't assume cloud parity. It resolves complexity in a shared risk context.

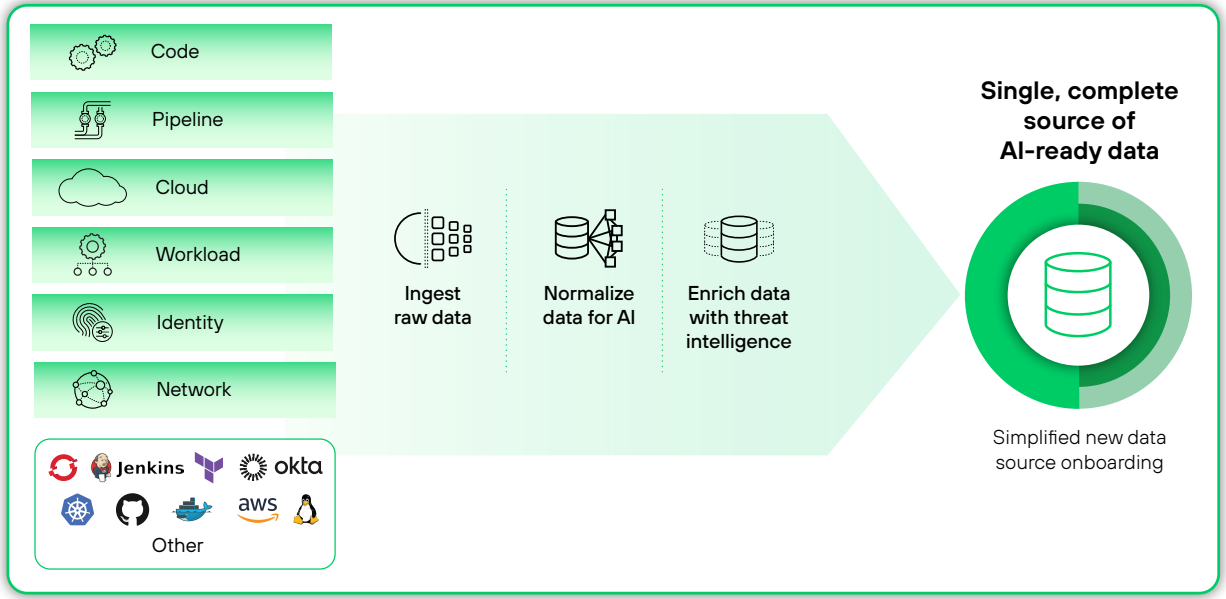


Figure 2. All data sources unified into one AI-ready stream

How does your platform use AI to improve security outcomes and reduce manual effort?

Most security AI expectations have gone underrealized because models lack the labeled data and scale needed for precision. Cortex Cloud, built into the Cortex platform, leverages over 10,000 detectors with more than 600 dedicated to cloud and aligned to the MITRE ATT&CK® framework. Analysts can instantly understand whether a signal reflects credential access, initial access, lateral movement, or exfiltration. Every mapped technique reflects active threat research from Unit 42®, allowing defenders to triage by attack stage, not just signal count.

Supervised Learning for Prevention-First Action

Using supervised learning based on labeled data drawn from full lifecycle telemetry allows the platform to understand both the existence of a vulnerability and its real-world exploitability and business impact. Static findings are correlated with runtime behavior, privilege escalation paths, and asset criticality to determine which risks require action.

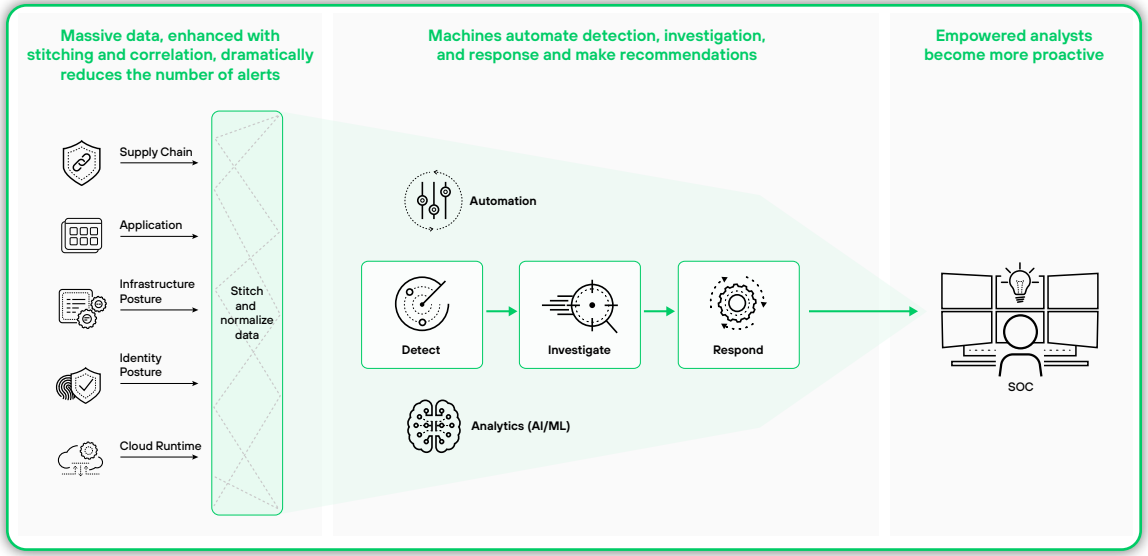


Figure 3. Unifying data sources enables AI-driven automation

Automation extends the Cortex Cloud prevention-first approach by enforcing guardrails throughout the software lifecycle and closing the loop between runtime detection and upstream controls.

When an exploitable issue is found in production, the platform can automatically update policies to prevent similar problems in future builds, ensuring issues are permanently fixed rather than repeatedly addressed. Additionally, the platform's embedded orchestration engine scales these actions across complex, multicloud environments, revoking risky permissions and rolling back vulnerable artifacts or triggering targeted playbooks based on context-driven criteria.

The result is a platform that both detects and responds faster, as well as continuously reduces tomorrow's attack surface.

How does your CNAPP support secure, role-based access and customizable dashboards without fragmenting visibility?

Most CNAPPs ask you to choose between customization and cohesion. Security engineers need runtime telemetry, posture context, and threat indicators. Meanwhile, developers want scoped views into their own pipelines and repository risk, and governance teams want to focus on compliance or audit posture. Addressing these needs often comes down to duplicating dashboards or reconfiguring permissions, perhaps running entirely separate components. Regardless of the method, neither workaround can surmount the problem with partial visibility and mismatched logic.

Role-Based Access on Unified Data

Cortex Cloud is built for cohesion. Where policies, issues, and cases exist as unified objects, they're consistently represented across every module. Views can differ by user, but data doesn't fork as it's pulled from and pushed to the platform's unified data lake.

Role-based access control (RBAC) and scope-based access control (SBAC) define what users see and interact with, down to the asset, cluster, or tag level. Security teams can scope dashboards, saved filters, and workflows by persona, business unit, or asset group.

What architectural choices simplify onboarding and accelerate time to value?

Security platforms can lose momentum in the first mile, slowed by infrastructure prerequisites, overpermissioned onboarding, or fractured module design. These early obstacles become delays that typically require help from internal cloud engineers or extended onboarding cycles. Cortex Cloud removes these hurdles.

Agentless-First Architecture

With Cortex Cloud, onboarding consolidates into a single flow. Once customers connect their cloud accounts, they immediately gain agentless visibility into cloud assets, vulnerabilities, Kubernetes clusters, datastores, and identities. The platform scans registries and workloads directly from the cloud with minimal permissions—just snapshot access, no VM provisioning or infrastructure deployment.

Taking an agentless approach eliminates the bottlenecks where security teams wait for infrastructure approval or extended cloud engineering cycles. Customers see posture visibility, configuration risks, and vulnerability exposure from day one.

Policy Abstraction Above Cloud Accounts

Policy enforcement is also decoupled from cloud-native constructs. Customers define guardrails and detection logic and apply them across assets or teams just once. The abstraction layer shortens the path from connection to control, giving security teams fast traction without burdening DevOps or cloud administrators.

Unlike legacy CNAPPs, where each module requires a separate RBAC setup and permission configuration, Cortex Cloud uses unified agent and identity management. Everything is natively integrated, involving no patchwork acquisitions or half-integrated tools.

If we're not ready to deploy everything at once, can your platform support staged adoption?

Traditional CNAPPs force teams to light up entire platforms before gaining meaningful insight, an approach that doesn't scale across all organizations. Cortex Cloud eliminates this constraint through flexible adoption paths that deliver immediate value regardless of starting point.

Multiple Entry Points with a Single Foundation

Customers can begin wherever their priorities dictate. Some start with code security to lock down their CI/CD pipelines. Others begin with posture management to assess configuration risks or runtime security to stop attacks in real time.

From the initial agentless foundation, organizations can add protection incrementally. They might deploy agents or Kubernetes connectors for deeper runtime visibility. The architecture delivers speed, independent of choice and without shortcuts. New capabilities require no reconfiguration or duplication of effort. They simply plug into the existing model to enrich the all-important big picture.

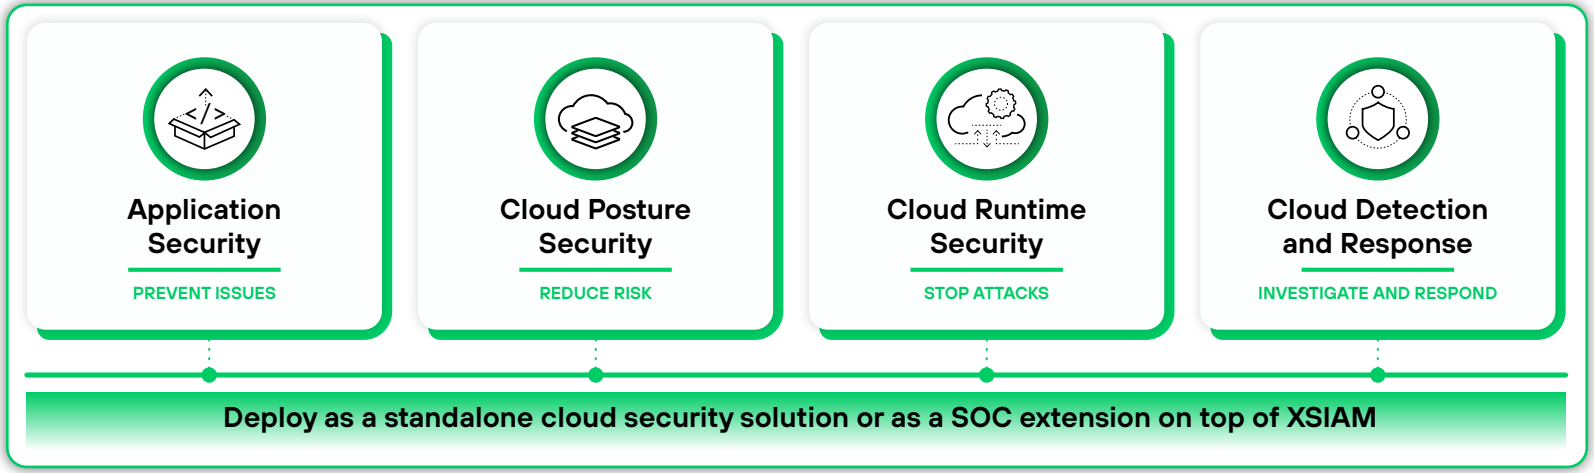


Figure 4. Unified coverage across application, posture, runtime, and response

2 APPLICATION SECURITY

Can you consolidate findings from existing AppSec tools into a single risk, policy, and automation engine?

Shift-left security was meant to reduce late-stage risk by catching issues in the pipeline where they're 10 times faster to fix. Without complete context, though, developers can't confidently ship code, and tools can't enforce intelligent guardrails. Cortex Cloud replaces reactive workflows with a prevention-first design built on unified data from code to cloud to SOC.

Prevention-First Integration

Cortex Cloud integrates directly into developer workflows, evaluating code at the point of change. Security checks run in GitHub, GitLab, IDEs, and registry promotions to flag issues in pull requests (PRs) and builds before they reach production. Developers see feedback inline—whether in PR comments or build results—without leaving their tools.

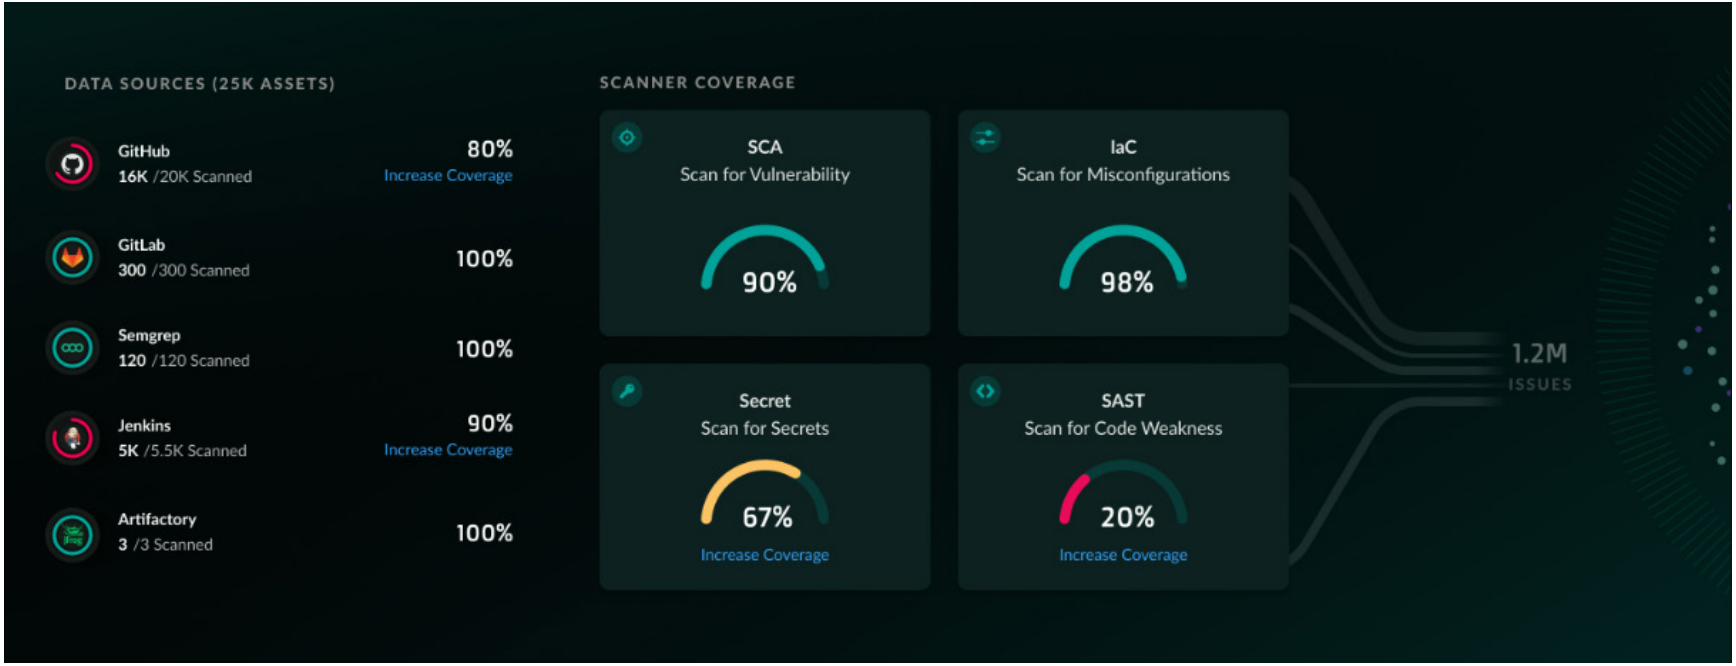


Figure 5. Visibility tracks coverage across repositories and development pipelines

At the same time, the platform ingests findings from industry-leading AppSec tools—through native integrations or a SARIF import—into a single risk, policy, and automation engine. Rather than treating each scanner's output as a separate stream, Cortex Cloud normalizes results into its object model, ensuring the same vulnerability or misconfiguration is tracked once across its lifecycle. Security teams gain visibility and control across heterogeneous tools, while developers continue working in the systems they already use.

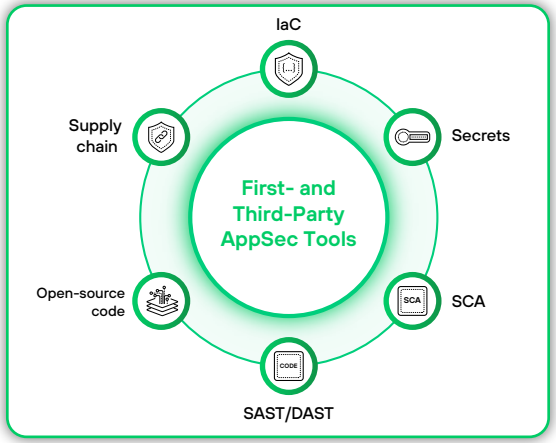


Figure 6. Integrating first- and third-party AppSec tools into one ecosystem

Can your CNAPP show the full chain from repository to pipeline, to image to runtime?

Many CNAPPs scan containers in registries or at runtime but lose the thread between how those images were built and where they now run. When a CVE appears in production, security teams struggle to trace it back to the repository, pipeline, or registry that introduced it, and their response becomes guesswork.

Cortex Cloud maintains that thread as a persistent object graph, tracking every image from commit to build to registry to runtime. Each image view links code provenance, metadata, and runtime state so analysts know that an image is vulnerable, as well as who introduced it, when it changed, and whether it's actively exposed. Risks detected upstream are checked against what's running, and runtime issues are traced back to the source for a permanent fix.

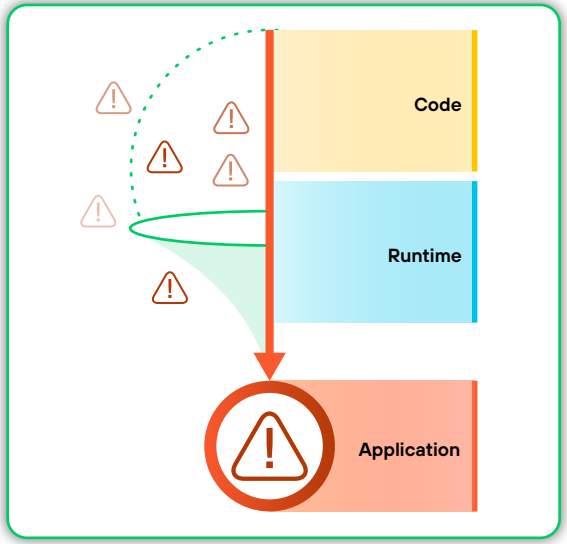


Figure 7. Context correlated across code, runtime, and application exposes critical risks

Risks detected upstream are checked against what's running, and runtime issues are traced back to the source for a permanent fix.

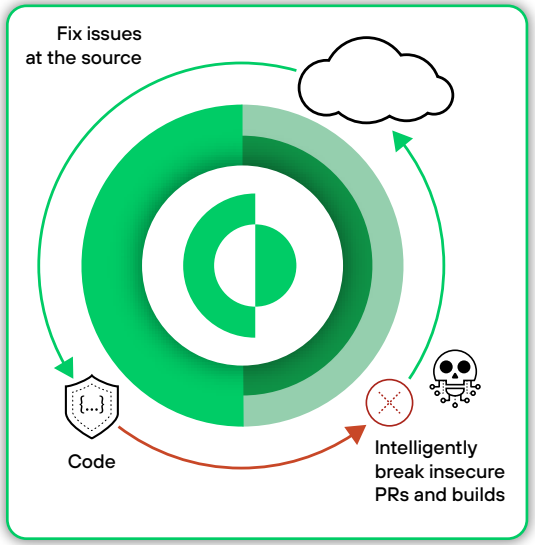


Figure 8. Intelligent guardrails block risky pull requests before deployment

What policy-as-code options exist for targeted blocking or reporting in PRs and builds?

Security breaks down when policies don't evolve with the environment. Static guardrails either slow developers with blanket restrictions or leave new projects outside enforcement until someone reconfigures.

Cortex Cloud expresses guardrails as code, defined in YAML or via API templates, and applied directly in CI/CD workflows. Pull requests can be blocked, builds failed, or registry gates enforced when high-risk conditions are detected. Lower-severity issues can be logged or pushed into issue tracking systems, giving security teams flexibility to match enforcement to context rather than enforcing a one-size-fits-all gate.

10X

Code risks take 10× longer to resolve in production.

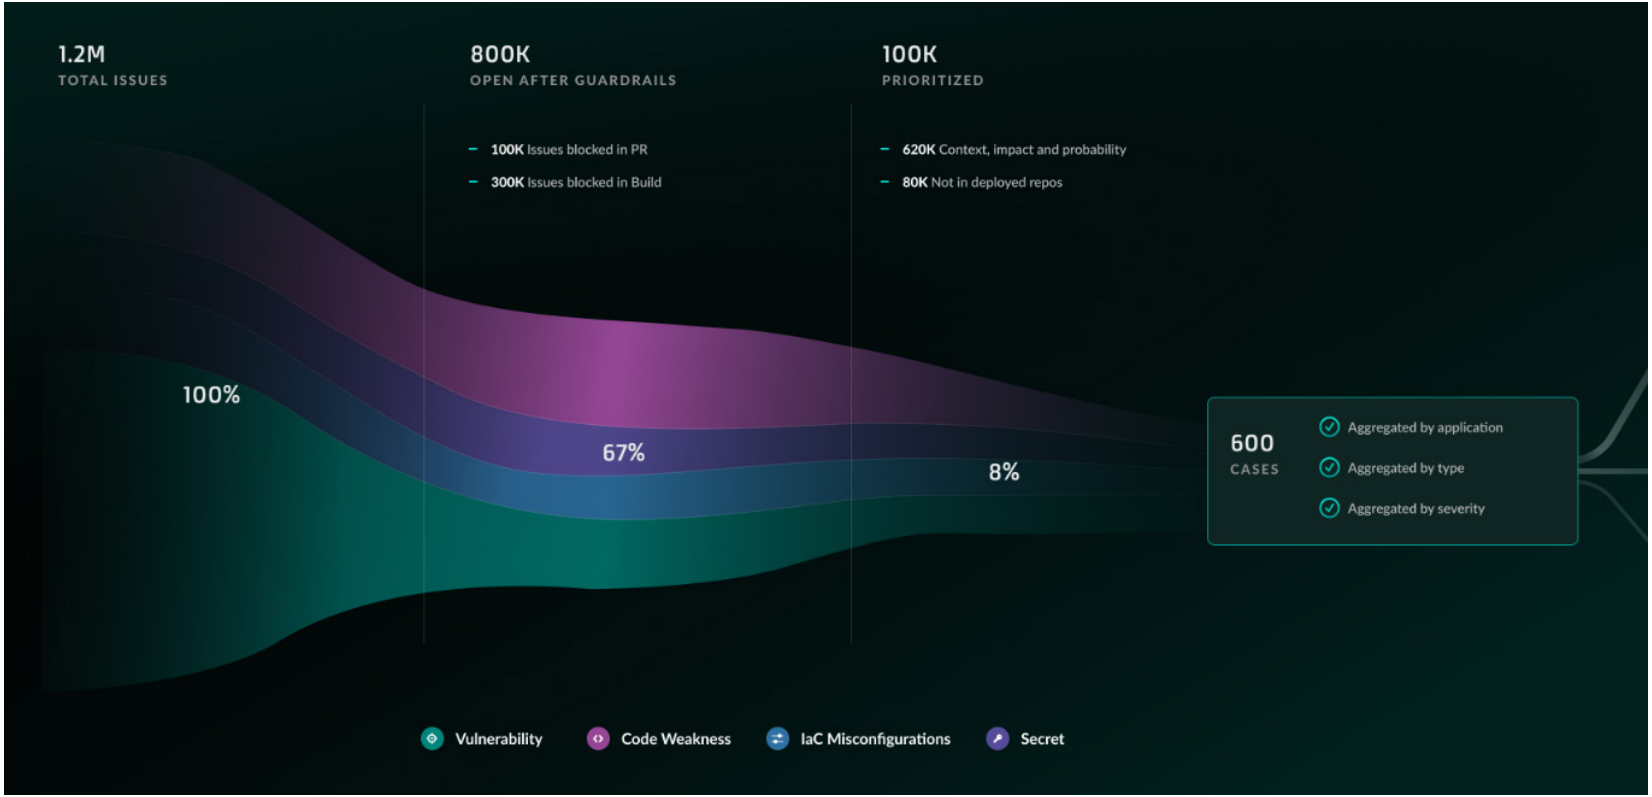


Figure 9. Guardrails and aggregated cases prioritize vulnerabilities by impact, type, and severity

Beyond enforcement, scopes adapt automatically as engineering scales. When a new repository, service, or team comes online, Cortex Cloud applies the appropriate guardrails without manual setup. Security teams define rules once—at the repository, pipeline, or organization level—and the platform ensures consistent coverage as projects multiply. Developers keep moving at velocity while security teams gain assurance that new code is governed from day one.

How does your CNAPP reduce friction between developers and security teams?

Because Cortex Cloud maps every issue to the repository, service, or workload that introduced it, ownership is never in doubt. Instead of security teams assigning tickets manually, issues flow automatically into Jira or ServiceNow for the responsible team. Developers don't need to chase context or switch consoles. They receive context-rich findings in the tools they already use, while security maintains centralized guardrails. The result is fewer handoffs, faster resolution, and collaboration that scales without slowing development.

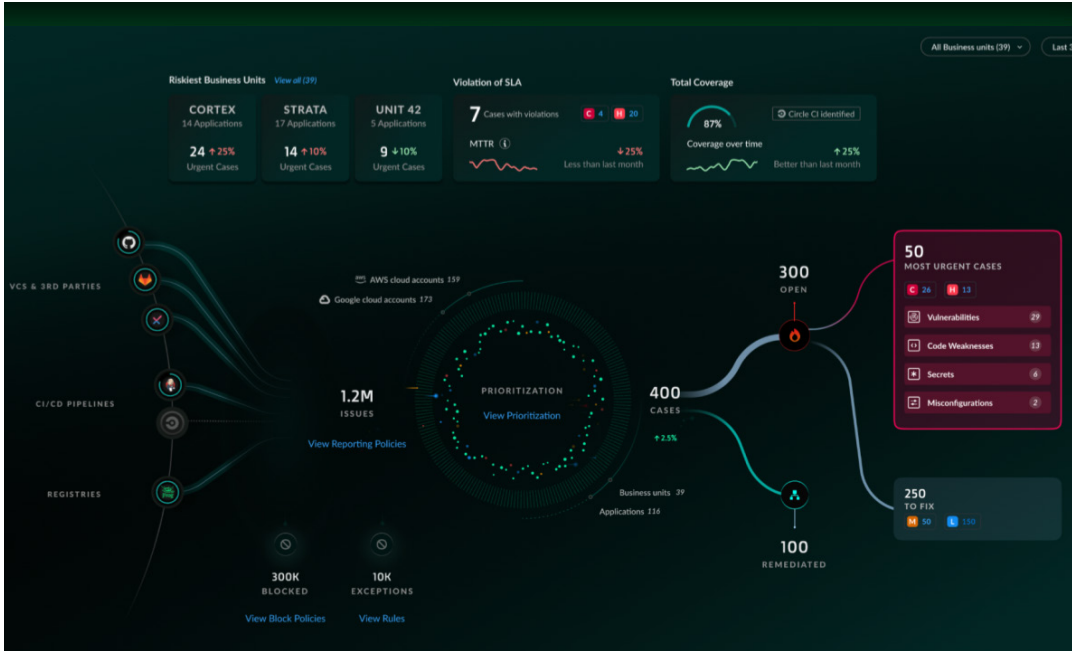


Figure 10. Visibility and automation expedite application security case management

3 CLOUD POSTURE AND IDENTITY SECURITY

How does your CNAPP help teams prioritize and remediate risks across cloud infrastructure?

Cloud posture produces overwhelming telemetry, but raw alerts don't drive action. Security teams need risk distilled into what matters, expressed in a way that enables coordinated remediation.

Cortex Cloud cuts through the noise in real time, promoting only findings that represent actual risk. Issues arrive with context already attached—identity reach, runtime activity, data exposure, and exploit paths. Instead of spinning cycles parsing static severity scores, teams get a ranked list of issues that require attention.

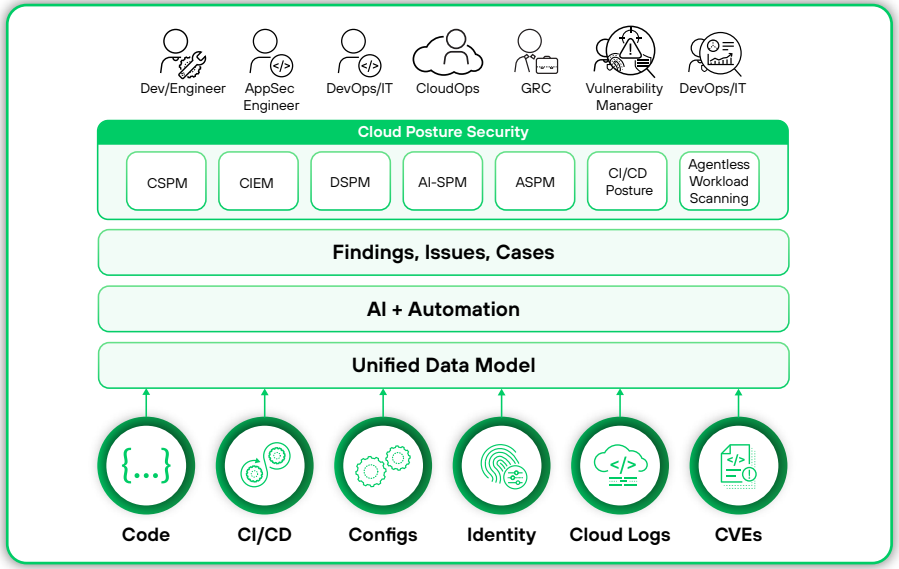


Figure 11. One data model powers AI-driven security from code to cloud to SOC

Filtered Findings, Prioritized Issues

Rather than alerting on everything that's discovered, the platform uses a structured model to classify and evaluate *findings*, which are simple facts about the environment that aren't necessarily good or bad. Only when a finding represents real risk does the platform escalate it to an issue.

Prioritization is dynamic. If a workload becomes internet-exposed or a dormant package executes, its *issue* moves to the top. But if compensating controls break the attack path, its weight falls. Triage decisions that once took hours are made instantly.

Issues are then rolled up into *cases*—not loose groupings, but attack-path narratives that connect weaknesses into breach risks. Each case ties together affected assets, identities, and permissions, as well as related artifacts so the chain of risk is clear end to end. A case becomes the operating unit for remediation, allowing teams to resolve an attack path in one motion instead of chasing alerts in fragments.

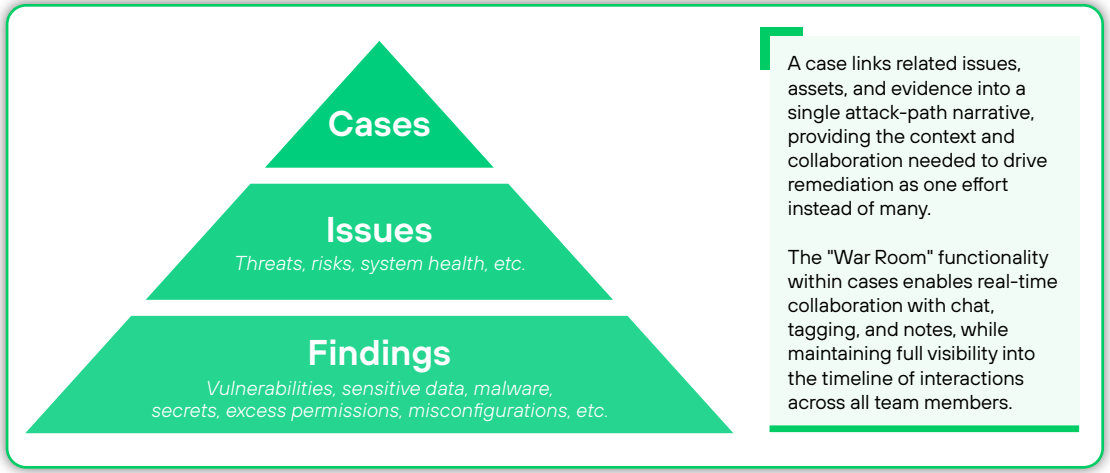


Figure 12. A structured hierarchy transforms raw findings into actionable security outcomes

Prevention at Scale

When team coordination is required, cases expand into a war room or investigation space that anchors all context, evidence, and ownership in one place. Teams can tag each other, comment, and automate handoffs inside the platform. Cases also track who's viewed or responded, reducing rework and ensuring accountability. Shared context is preserved, as it's built into the system.

Remediation is both guided and scalable. Cases propose targeted changes, tightening a role, closing a storage exposure, and patching a vulnerable image. Those changes can then be applied once and propagated everywhere the same pattern exists. Fixes land where they're most durable, upstream in code through policy enforcement and downstream in runtime through automated orchestration. Each fix is validated and locked in by guardrails, ensuring the same issue doesn't reappear.

By aligning issues, cases, and remediation in one architecture, Cortex Cloud transforms posture management into prevention at scale.

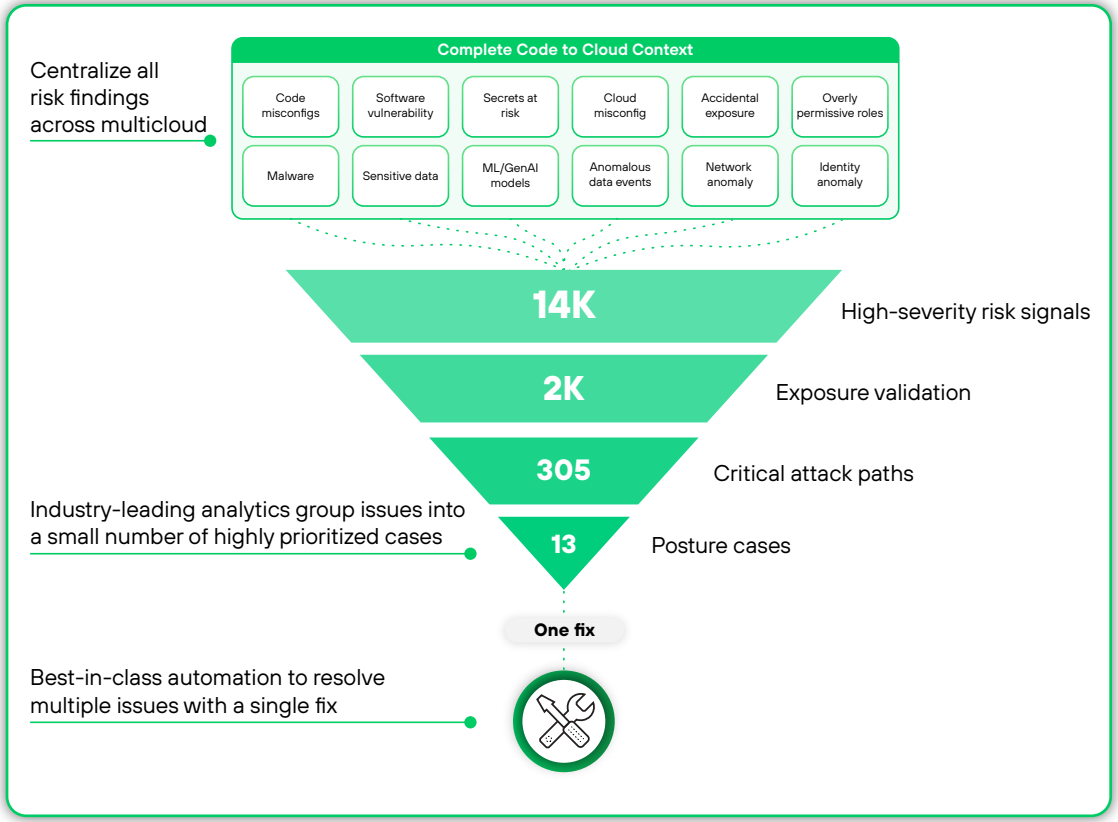


Figure 13. Analytics distills thousands of risks into prioritized, fixable cases

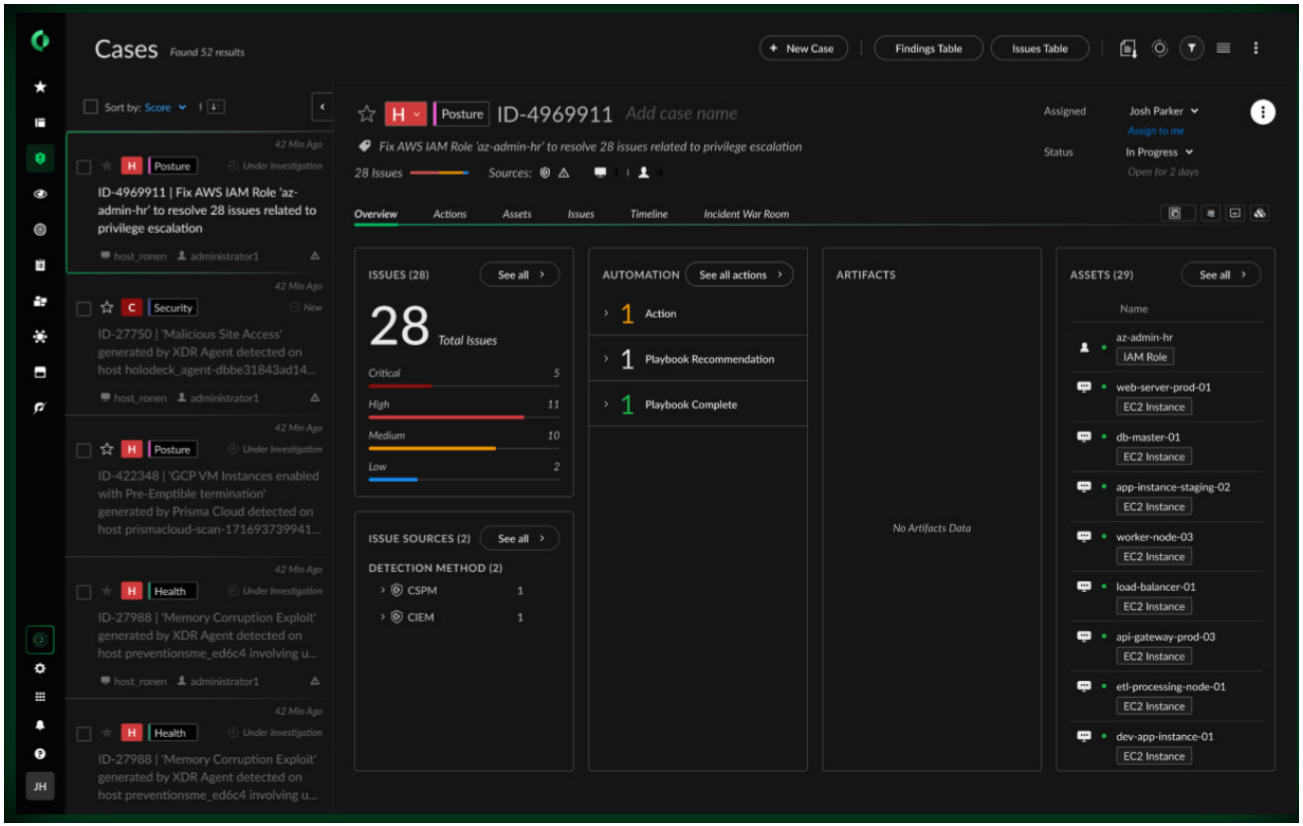


Figure 14. The cases framework guides one action to remediate issues across assets

Does your CNAPP preserve context and evidence from short-lived cloud assets to support forensic investigations?

Containers launch and terminate in seconds. Functions run on demand. Traditional forensic methods built for long-lived systems can't capture behavior from assets that no longer exist.

Cortex Cloud retains investigation depth across ephemeral workloads by streaming runtime telemetry to the unified data backend while the asset is live. When agents are deployed, they continuously monitor process activity, file access, network behavior, and identity usage. The data is sent off-host immediately and enriched with posture and identity context. Findings and signals tied to each workload persist in the platform's asset graph after the underlying container, function, or virtual machine terminates.

Cases that involve short-lived assets remain fully actionable. Analysts can see when a container ran, what policy violations it triggered, what roles it assumed, and how it interacted with other resources. Its history becomes part of the case timeline, alongside static risks, entitlement data, and cloud configuration insights. Even when full runtime observability is limited, posture and identity telemetry fill in key context across the attack lifecycle.

Cortex Cloud doesn't rely on logs alone. It builds a structured record of workload behavior tied to persistent entities in the platform. Analysts can investigate threats involving infrastructure that no longer exists and still reconstruct what matters most—what the asset did, what it exposed, and how to prevent it from happening again.

How does your platform support compliance, audit readiness, and data residency across global regions?

Many CNAPPs leave teams scrambling to meet regulatory requirements with findings scattered across modules, access controls lacking audit trails, and data storage spanning regions without residency guarantees. Cortex Cloud addresses compliance through platform-wide controls and traceability, as well as regional isolation built into the foundation.

Embedded Compliance Controls

Security controls are embedded in every stage of the software development lifecycle. Vulnerability scans, infrastructure-as-code validation, and production access controls are enforced continuously with actions logged for audit review. Just-in-time access is tightly scoped, time-bound, and fully audited.

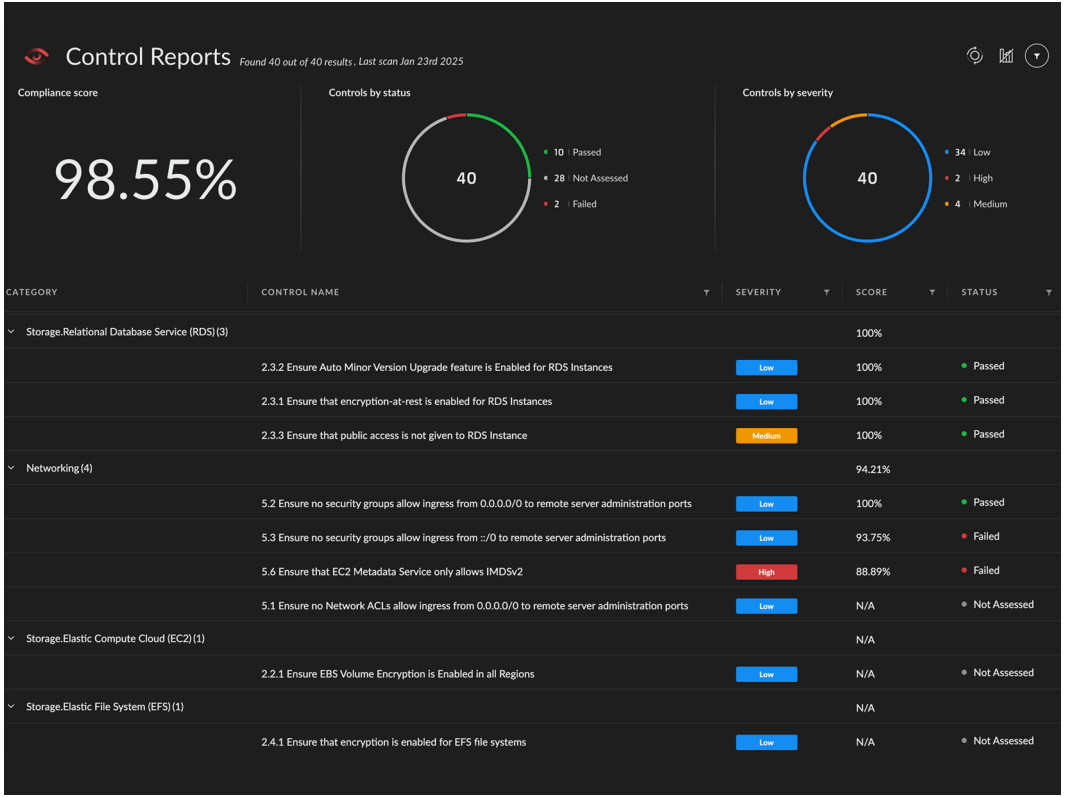


Figure 15. Clear visibility into the control status strengthens the cloud governance posture

The platform maintains comprehensive audit trails across all security activities, from policy enforcement to incident response. Teams can query activity and demonstrate enforcement directly within the platform without exporting logs or building custom compliance reports.

Standards Certification and Framework Alignment

Cortex Cloud is certified against leading security and compliance standards, such as SOC 2 Type 2 and ISO 27001. Customers can align their controls and reporting to industry benchmarks without customizing tools or rebuilding visibility.

Prebuilt compliance templates with predefined reports and configurations to help automate compliance reporting for regulatory frameworks include:

- CIS
- SOC 2
- ISO
- NIST
- HIPAA
- HITRUST CSF
- SOX
- PCI DSS
- GDPR
- CCPA
- FedRAMP
- OWASP
- MITRE ATT&CK

The platform's unified data model enables consistent compliance reporting across all security domains—posture, identity, runtime, and application security—eliminating the gaps that arise when compliance data is fragmented across tools.

Data Residency and Sovereignty

Each customer operates in a segregated tenant within their selected cloud region. All telemetry, findings, metadata, and backups are stored in that region and never leave it. Even AI and ML processes execute locally to preserve data sovereignty and support region-specific compliance mandates.

Regional isolation ensures organizations can meet data residency requirements for regulations like GDPR while maintaining full platform functionality. Compliance becomes an architectural guarantee rather than a configuration challenge.

4 RUNTIME PROTECTION AND RESPONSE

What runtime protection options are available, and how do they cover modern compute types?

Cortex Cloud provides flexible runtime protection through multiple deployment models that balance security coverage with operational requirements.

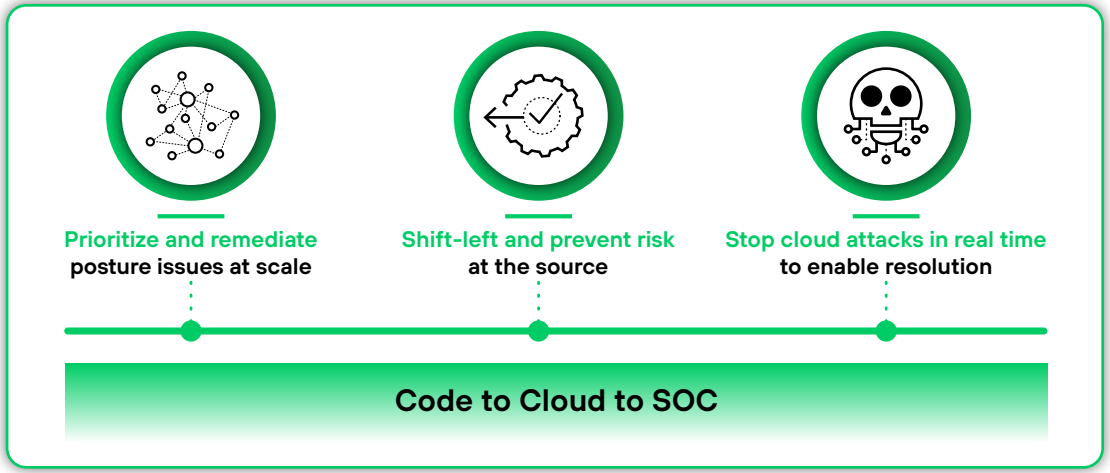


Figure 16. Continuous monitoring enables immediate defense against active threats

Agentless Foundation

Agentless protection provides immediate visibility without any deployment overhead. The platform snapshots images and workloads for analysis, inspects logs, and pulls configuration data directly via cloud APIs. Agentless scanning delivers fast, comprehensive coverage for vulnerability assessment, configuration analysis, and compliance checking without requiring infrastructure changes or performance considerations.

While agentless protection excels at discovery and posture assessment, it can't provide real-time threat detection or active blocking capabilities because it lacks deep integration with running workloads.

Full Runtime Agent Protection

For comprehensive runtime security, Cortex Cloud deploys the Cortex XDR® agent—a single-agent framework that provides complete protection across virtual machines (Linux) and containers. The agent delivers real-time threat detection through event collection, behavioral threat protection, and antimalware capabilities.

The agent analyzes files locally and communicates with WildFire® for advanced threat intelligence, enabling detection and blocking of zero-day attacks, malware, and suspicious behaviors. Customers can configure agents in detect-only mode for monitoring or prevent mode for active threat blocking, providing full control over the security posture.

Kubernetes Protection

In Kubernetes environments, Cortex Cloud deploys the agent as a DaemonSet, running a privileged pod on each node. From that position, it monitors activity across other pods and enriches runtime telemetry with Kubernetes context, such as pod name, namespace, and container image, which ensures visibility extends beyond the host to the cluster and container levels. Cortex also supports agentless auditing of Kubernetes control-plane activity, complementing the node-level agent to provide broader cluster awareness.

Serverless and CaaS Protection

For container-as-a-service environments and serverless functions, Cortex Cloud provides specialized agents with more limited capabilities due to platform constraints. These agents focus on blocklist enforcement and basic antimalware protection rather than full behavioral analysis.

Essential Security Mode

For environments where resource consumption is a primary concern, the platform offers an Essential Security mode. The configuration reduces capabilities by disabling event collection and behavioral threat protection while maintaining antimalware and antiexploit functionality.

Flexible Deployment Strategy

Customers can mix deployment models based on workload criticality, where critical workloads run with full agents for maximum protection, and less critical environments use agentless scanning for speed and simplicity. The platform also supports policy-driven automation, allowing workloads to move between protection modes as conditions change. Because policies apply seamlessly at the agent layer, there's no need to reinstall, restart, or redeploy—coverage adapts in real time without operational overhead.

How does your platform detect evasive or unknown threats in real time?

Traditional runtime protection forces teams to choose between signature-based detection that attackers easily evade or anomaly detection without context that flags every deviation from baseline.

Cortex Cloud uses a hybrid detection strategy that combines local prevention, behavioral rule engines, and machine learning trained on labeled threat intelligence.

Multilayered Detection Architecture

When unknown files execute, the runtime agent uses local static analysis and ML-based verdicting to determine maliciousness. If local analysis proves inconclusive, the platform escalates to WildFire for detonation and confirmation. The tiered model ensures both speed and precision without relying solely on external analysis.

The platform leverages the Palo Alto Networks substantial data advantage—models trained and validated across diverse environments spanning financial, industrial, and tech sectors in over 130 countries. New customers benefit from detection capabilities already tuned at massive scale rather than serving as early test cases.

How Unit 42 Can Help

Unit 42, the Palo Alto Networks dedicated threat intelligence and incident response team, works with customers to provide a range of cybersecurity services, from proactive measures to emergency support.

If you're experiencing a security incident, the Unit 42 Incident Response team is available 24/7 to help your organization investigate, contain, and recover from the attack. Customers can also purchase an incident response retainer, which ensures faster access to the team's expertise.

In addition to incident response, Unit 42 offers managed services, like Managed Detection and Response, and cybersecurity consulting to help organizations improve their security posture and build more resilient defenses.

For immediate help with a security incident, contact the Unit 42 Incident Response hotline.

Behavioral Threat Protection Engine

Beyond anomaly detection, Cortex Cloud applies thousands of curated rules developed and validated by Unit 42 researchers. Each rule defines specific malicious patterns rather than simple deviations from normal behavior.

For example, uploading a file to a server isn't malicious alone, but when it follows unauthorized access and other suspicious activity, the behavioral engine identifies the complete pattern as malicious. Rule-based correlation enables detection and blocking of zero-day attacks by recognizing attack sequences rather than individual indicators.

MITRE ATT&CK Mapping and Context

Detection extends across the attack lifecycle, with agent-mapping alerts to MITRE ATT&CK tactics and techniques. When threats surface, they carry rich annotations explaining the attack phase and severity, enabling analysts to respond based on attacker intent.

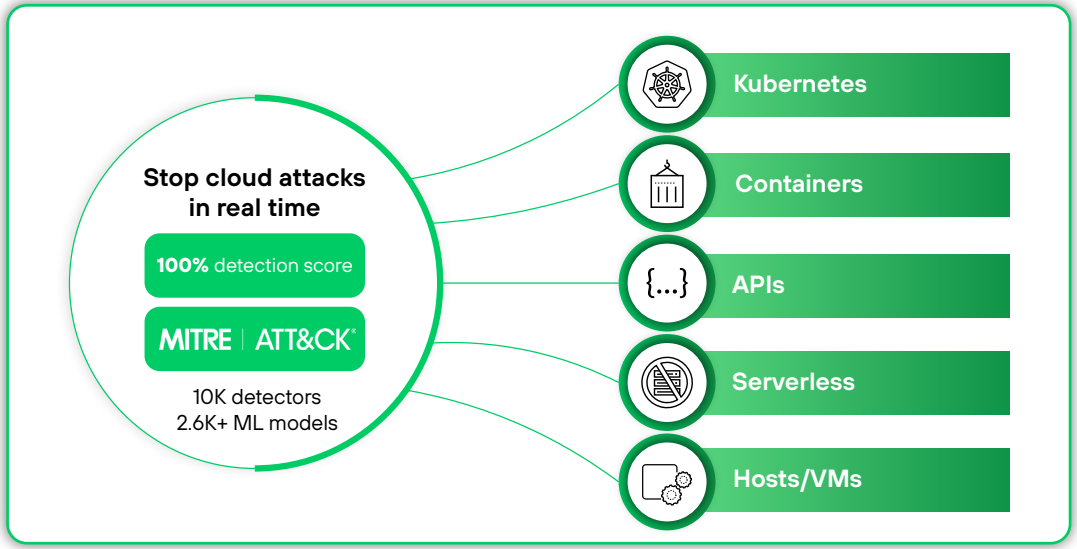


Figure 17. Aligning defenses with the MITRE framework stops cloud attacks across every workload

The Cortex platform has achieved 100% technique-level detection in MITRE ATT&CK Evaluations with high accuracy scores, demonstrating the platform's ability to detect sophisticated attack patterns while maintaining low false positive rates. The approach delivers runtime protection that understands context, recognizes behavior, and adapts as attackers evolve, proving that unknown doesn't mean undetectable.

MITRE ATT&CK Validation

The Cortex Cloud runtime protection capabilities are powered by the same Cortex XDR agent that has achieved industry-leading results in MITRE ATT&CK Enterprise Evaluations, Round 6—the cybersecurity industry's most rigorous endpoint security testing—for five years running.

Cortex XDR excelled in both detection and prevention scenarios of the evaluation:

- **Detection coverage:** First vendor to achieve 100% technique-level detection with no configuration changes or delays.
- **Prevention performance:** Highest prevention rate among all vendors with zero false positives.
- **Platform coverage:** 100% detection across Windows, Linux, and macOS attack surfaces.

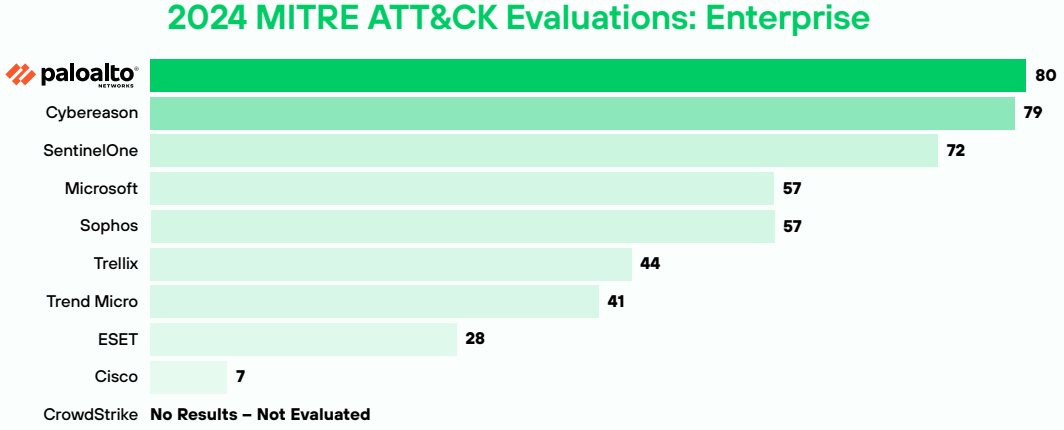


Figure 18. Palo Alto Networks excels in the recent MITRE ATT&CK Evaluation

MITRE ATT&CK Evaluations test real-world attack scenarios, including ransomware campaigns and advanced persistent threats. Technique-level detections provide security analysts with precise information about attack progression, enabling faster response and threat containment.

The same behavioral analysis engines and threat detection capabilities validated through MITRE testing integrate directly into Cortex Cloud, ensuring runtime protection maintains enterprise-grade effectiveness while operating within the broader cloud security context.

Read the results of MITRE ATT&CK Enterprise Evaluations, Round 6.

5 SOC INTEGRATION AND OPERATIONAL SCALE

How does your CNAPP turn code-to-cloud risk and SOC operations into one continuous motion across every security team?

Security outcomes falter when cloud risk lives outside the record of work. Our platform centers operations on a single data plane and an incident-grade object so every signal—from code, cloud, identity, and runtime—resolves to the same assets and principals, with ownership and next actions attached. Investigations run on that object, and closure feeds prevention so context doesn't fragment and work doesn't duplicate.

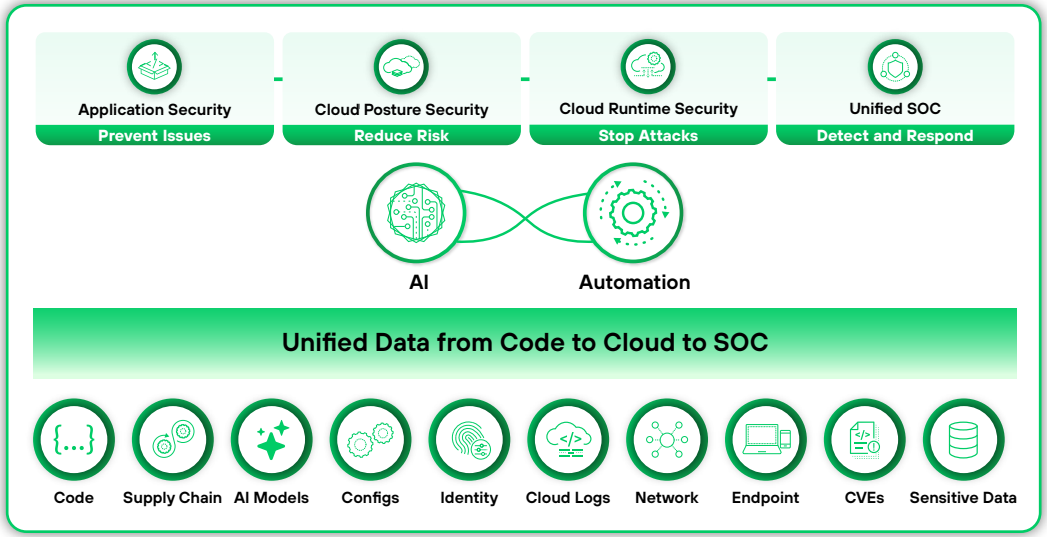


Figure 19. Incident-grade objects align signals, context, ownership, and prevention

How It Works End to End

One normalized model binds durable IDs to assets and identities across environments, preserving risk lineage from change to workload to account. A shared state machine carries the lifecycle from open to prevented with bidirectional sync so comments, dispositions, and playbook outcomes persist everywhere without retyping or drift. With Cortex Cloud, teams investigate faster because correlation is precise and deduplication happens by entity, not string match.

What Every Team Gains

SOC and incident response work a single incident narrative with evidence bundles and ready-to-run actions. Threat intelligence sees lineage inside the investigation and can target the right control points. AppSec and CloudSec get ownership mapped to the teams that can fix issues, and closure updates policies and guardrails so problems don't recur. Architecture and engineering operate one schema and method on record, not ad hoc exports.

Vulnerability management and IAM benefit from explicit entity resolution so duplicates collapse and principals remain consistent across systems. Governance, risk, and compliance (GRC) and data security inherit full audit trails and regional guarantees without rebuilding evidence.

Why Cortex Cloud Is Different

Most platforms stop at enrichment, moving rows of alerts into another console. Cortex Cloud operationalizes those signals as native incidents with lineage, ownership, and actions—no one-way webhooks, no duplicate state machines, no "open elsewhere" detours, and no brittle correlation that breaks on ID rotation. The SOC works from a complete cloud context, and prevention improves with every investigation.

Proof Buyers Can Run

A high-risk cloud finding appears once as a deduplicated incident with an associated playbook. When the case is closed, the posture updates automatically so the same issue doesn't resurface. The same human or machine principal resolves consistently across code, cloud, and runtime, and promotion into SOC workflows meets the agreed SLO. Results are measurable and reproducible, with outcomes attributed to the integration rather than lost in operational drift.

Bottom Line

One data plane, one incident object, one workflow. Every team—SOC, threat intelligence, incident response, GRC, AppSec, CloudSec, security architecture and engineering, vulnerability management, IAM, and data security—works from the same record of truth, turning fragmented steps into a single motion from finding to fix to prevention.

What cost, complexity, or staffing efficiencies can we anticipate if we go with your platform?

When each component of your tech stack adds its own interface, data model, permission schema, and operational overhead, security teams spend more time reconciling systems than reducing risk. CloudSec and DevOps field duplicate findings from disconnected scanners, while SOCs chase alerts that never connect to a larger narrative. Training multiplies, ownership blurs, and costs climb.

Cortex Cloud eliminates waste. One platform can serve as a CNAPP, a SIEM connector, an EDR console, an AppSec orchestrator, a runtime enforcer, and a SOAR engine. Teams investigate in one place, speak a common language, and act without toggling systems.

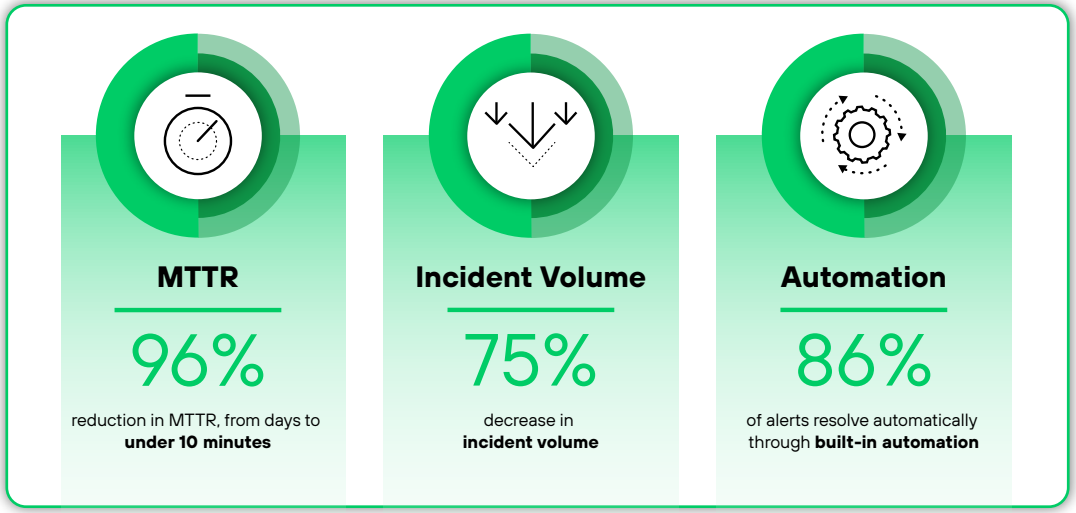


Figure 20. Aligning defenses with the MITRE framework stops cloud attacks across every workload

Measurable Performance Improvements

Organizations using the Cortex platform see MTTR drop from days to under 10 minutes—a 96% reduction. Incident volume decreases by 75%, while 86% of alerts resolve automatically through built-in automation.

The case-based model eliminates alert chasing. Analysts work with complete narratives that dramatically accelerate detection and response timelines.

Consolidation Benefits

Customers avoid the overhead of managing multiple vendors, including licensing negotiations, integration projects, upgrade cycles, and training across disparate tools. Instead of managing five vendors, teams operate one unified platform.

Staffing and Training Optimization

Because teams don't need to master and maintain multiple consoles, training requirements decrease while operational breadth increases, enabling security programs to gain scale without proportional headcount growth. Plus, automation capabilities offload manual work with over 1,000 prebuilt playbooks that free analysts to focus on strategic security initiatives.

Advanced Attack Detection

Cost reduction represents only part of the value equation. Advanced attacks increasingly span endpoint and cloud environments. Fragmented tooling simply can't detect multidomain threats. Cortex Cloud correlates endpoint movement, cloud access, and command execution in one place, delivering detection capabilities that are impossible when piecing together point solutions.

The unified platform frees teams from platform management burden, enabling focus on driving risk reduction across every layer of the modern attack surface.

Ready to Imagine the Possibilities?

Contact us to discover how a unified security strategy can empower your organization to outpace modern threats—or schedule a demo of Cortex Cloud today. Explore the power of platformization.

About Cortex Cloud

Cortex Cloud, the next generation of Prisma® Cloud, merges best-in-class CDR with industry-leading CNAPP for real-time cloud security. Harness the power of AI and automation to prioritize risks with runtime context, enable remediation at scale, and stop attacks as they occur. Bring together your cloud and SOC on the unified Cortex platform to transform end-to-end operations. Experience the future of real-time cloud security at https://www.paloaltonetworks.com/cortex/cloud.

paloalto® NETWORKS

3000 Tannery Way
Santa Clara, CA 95054

Main: +1.408.753.4000
Sales: +1.866.320.4788
Support: +1.866.898.9087

www.paloaltonetworks.com

© 2025 Palo Alto Networks, Inc. A list of our trademarks in the United States and other jurisdictions can be found at https://www.paloaltonetworks.com/company/trademarks.html. All other marks mentioned herein may be trademarks of their respective companies.

cortex_eb_the-cnapp-buyers-guide_091725