5 Requirements to Protect AI and Cloud Deployments

Actionable Guidance for Managing Risk

paloalto® NETWORKS | CORTEX® CLOUD

AI Is Driving the Modern Cloud

Cloud and AI adoption have transformed the way organizations build, deploy, and run their business-critical applications.

With more than 90% of organizations in the cloud embracing artificial intelligence (AI)-assisted coding,[1] rapid adoption of new technologies introduces new potential for businesses but also increases complexity for security teams.

Rapid Innovation from Cloud Providers

Organizations rapidly adopt data, AI, and ephemeral compute technologies across multicloud environments.

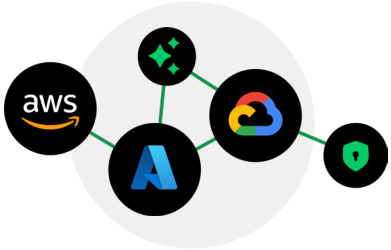

Skyrocketing Deployments

Cloud providers introduce new APIs and cloud services on a weekly basis.

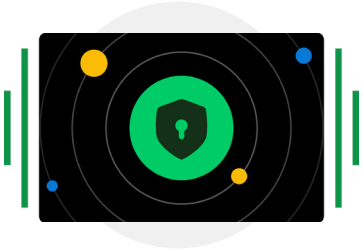

More Dangerous Attacks

AI enables threat actors to craft higher volume and more sophisticated attacks to gain access to cloud environments.

This e-book provides the guidance you need to ensure the best possible security outcomes in your multicloud environment—despite the challenges of the modern security landscape.

1. *The State of Cloud-Native Security Report*, Palo Alto Networks, March 7, 2025.

Tools sprawl creates cloud security challenges

On average, organizations adopt

16+ tools

for cloud security[2]

Cloud Security Challenges

Siloed context

Increasing cloud adoption drives the need for new security requirements and the adoption of additional tools such as:

- Cloud security posture management (CSPM)
- Cloud workload protection platform (CWPP)
- Data security posture management (DSPM)
- AI security posture management (AI-SPM)
- Cloud detection and response (CDR)

The result? Context is scattered across tools, and constant console switching bogs down security teams. Plus, it's almost impossible to achieve a complete understanding of the overall cloud estate.

Poor prioritization

With cloud security, more tools don't lead to improved visibility and reduced risk.

In fact, 90% of cloud-native organizations agree that the point tools they use create blind spots that affect their ability to prioritize risk and prevent threats.[3]

The simple reality is that disparate tools provide too many alerts and not enough visibility into overall risk. A vulnerability management tool, for example, identifies virtual machines with a critical vulnerability, but it doesn't recognize which machines have internet exposures or access permissions to crown jewel assets. To understand the data and prioritize risk, security teams must manually correlate findings that span multiple consoles.

2. *The State of Cloud-Native Security Report.*
3. Ibid.

Slow response times

Cloud security isn't just about visibility; teams must also take action to stop risks and threats. But remediation is difficult when siloed data, poor prioritization, and manual efforts stand in the way.

When organizations struggle to minimize their meantime to remediate critical issues, they're left vulnerable to threat actor scanning and exploits.

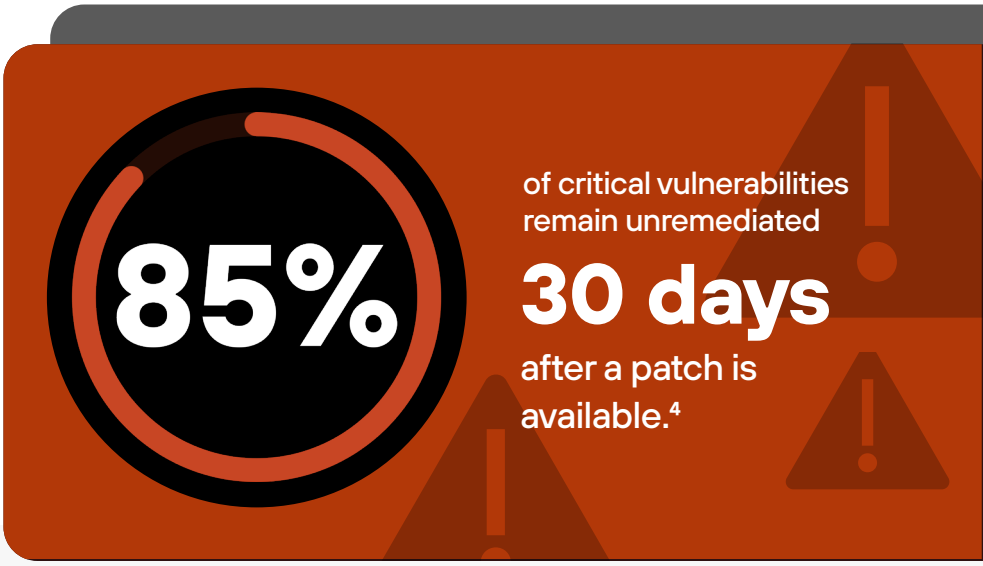

Friction points across internal teams

The modern cloud requires collaboration across different teams—from code to cloud to SOC—each with unique roles in securing the cloud environment. Application security teams monitor code, pipelines, and applications. Cloud security engineers safeguard infrastructure and workloads, extending protection to cloud-hosted data and AI models. SOCs thwart attacks and handle cloud incidents.

Developers

Build applications in the cloud

Application Security

Secure the development environment

Cloud Security

Secure the cloud environment

SOC

Detect and respond to attacks across the enterprise, including the cloud

As critical as all of these teams are to an effective security strategy, they operate with distinct tools and datasets, which complicate collaboration because they aim to secure the cloud comprehensively.

Security duties fall on security functions while development teams control cloud resources and settings. This split demands that security teams and developers work together to protect the cloud. The traditional security practice of handing off long lists of misconfigurations or vulnerabilities to developers without clear prioritization doesn't work. Among cloud-native organizations, 84% agree that security processes cause delays to developer project timelines.[5] Without a targeted architecture, security hinders developer productivity.

4. *Data Breach Investigation Report*, Verizon Business, 2024.
5. *The State of Cloud-Native Security Report*.

Five Essentials to Simplify Risk Management

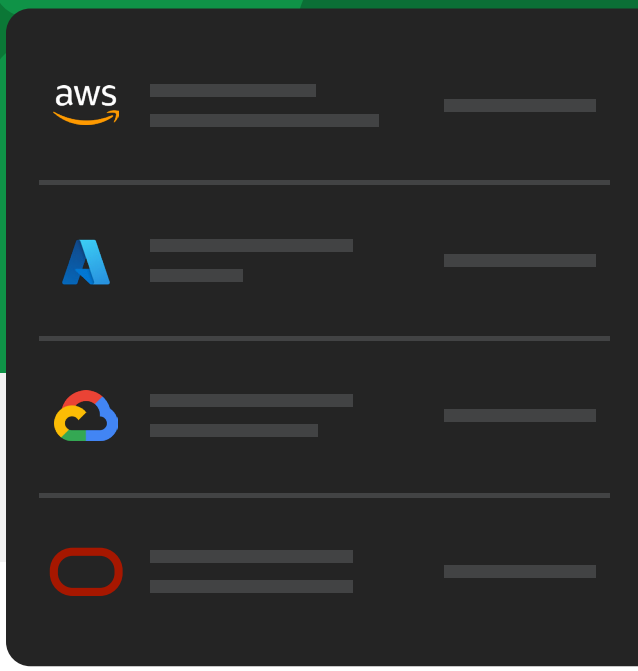

Ensure the best possible security outcomes by making it easier to manage risk.
Follow these steps to get started.

#1 Unify Data

Securing a multicloud environment starts by gaining visibility into all cloud resources, configurations, and activities. Then, it entails establishing a comprehensive overview for robust cloud security processes. Here's what to focus on for effective visibility:

Cloud infrastructure: Whether using a single provider or planning for multicloud adoption, design an architecture that spans major clouds like AWS, Microsoft Azure, Google Cloud, and Oracle Cloud.

Compute architectures: Understand the full scope of your compute environments, from virtual machines (Linux and Windows) to containers, Kubernetes, and serverless deployments. Be sure to include their configurations and content, like host OS, packages, data, and secrets.

Identity and permissions: Cloud providers offer complex identity and access management (IAM) frameworks. Aim for a solution that automatically detects and maps permissions, providing a clear picture of access rights across all assets.

Data: With cloud adoption, data usage has surged, often including personal identifiable information (PII) and other sensitive data in services like Amazon S3 or Azure Blob Storage. Therefore, scan and classify these datastores to pinpoint where sensitive information resides.

AI: Because of the rapid expansion of AI in the cloud, particularly for training large language models (LLMs), understand how and where AI is applied within your organization.

Deployment architecture: Traditional security systems with inline tools, like proxies and network scanners, are often excessive for mere visibility. Modern multicloud security favors an agentless approach, enhancing visibility without needing installed agents.

Beyond broad visibility, unifying data across cloud components—including network, identities, workloads, data, and AI—into a centralized data lake is key. This integration powers AI and automation-driven cloud security, while improving outcomes.

#2 Detect and Prioritize Risk

Cloud dynamics often lead to an abundance of misconfigurations, but not all security issues pose a real risk. For instance, a misconfigured, unused Amazon VPC is less critical than one that contains essential resources. Teams should prioritize addressing their most significant threats to prevent attackers from exploiting real vulnerabilities. This underscores the importance of risk prioritization.

Organizations should brace for cloud misconfigurations, compliance violations, vulnerabilities, overprivileged access, sensitive data exposures, and AI model risks. However, detecting these risk signals individually is insufficient. Therefore, integrate data across systems to achieve effective cloud security. Automated correlation helps minimize alerts and spot major issues.

Public Exposures

Resources exposed to the internet are prime targets for attacks. Ensure you have effective tools that can stitch together the entire network path—including access rules—to identify internet-exposed virtual machines and containers. Integrate identity and access policies with sensitive data context to highlight public data exposures as well.

Attack Path Detection

Not all misconfigurations are critical, but your teams must identify those that facilitate attack paths or lateral movements to vital resources. For example, a virtual machine that combines a critical vulnerability, internet exposure, and access to sensitive data poses a high risk.

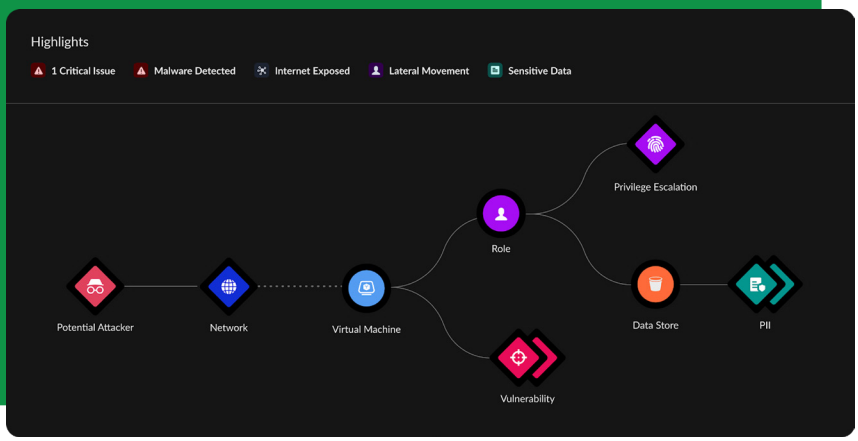

Alert Consolidation to Work Smarter

Cloud environments present numerous security issues and attack paths, often stemming from the same root cause. For example, a single AWS security group might account for numerous network exposures, or a single admin role could facilitate multiple lateral movement paths. Consolidating issues into fewer, contextualized alerts based on a common root cause is a more scalable approach than sifting through individual risks one at a time. Using AI-driven analytics to correlate data effectively enhances detection and prioritizes risks, compared to static-based detections found in manually configured policies.

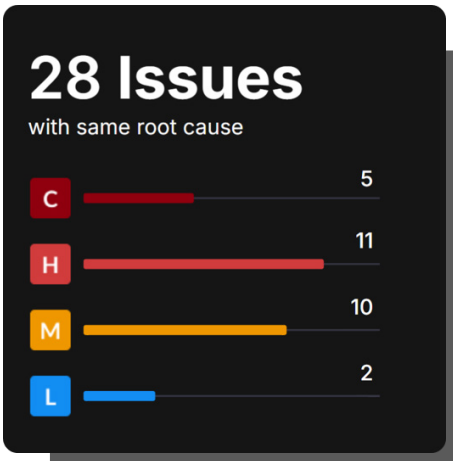

#3 Remediate Risks Efficiently

Maximizing your cloud posture goes beyond visibility and risk prioritization. It requires proactive risk elimination through effective remediation workflows. To avoid inefficiencies in the process, organizations must first understand the teams involved in remediation.

Cross-Functional Collaboration

Cloud security demands teamwork. Security teams pinpoint critical issues but often lack the authority to alter cloud configurations directly. That's because they don't understand the impact of cloud configuration changes. An improper change without adequate review, testing, and approval can disrupt applications or operations. Typically, the developers, DevOps, app owners, and cloud operators responsible for resolving these issues view security measures as productivity obstacles. To minimize risk without sacrificing agility, design security programs with developer considerations at the forefront.

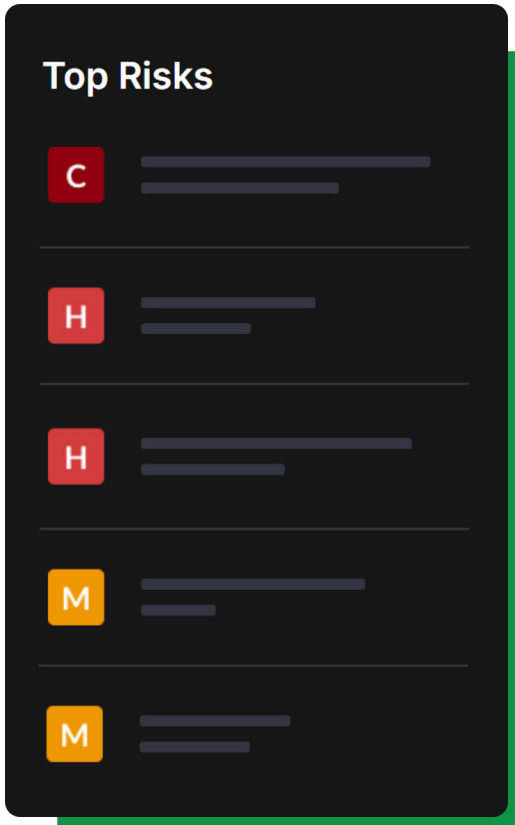

Fewer Issues and More Context for Developers

Developers primarily focus on creating new features or resolving bugs. They have limited capacity for addressing security issues so simply listing vulnerabilities for them to fix is ineffective. Negotiations between your development and security teams over fixes and timelines can benefit from providing developers with full risk context, like attack path analysis, to support remediation priorities. It's also helpful to provide a prioritized list of critical issues. Doing so enables your developers to choose fixes based on risk impact rather than simplicity.

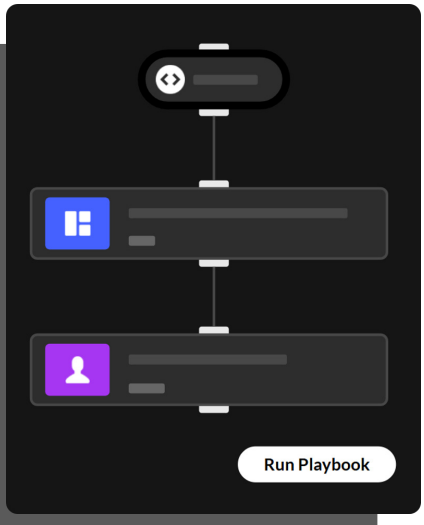

Efficient Remediation

It's unsustainable to address thousands of misconfigurations one at a time. Reduce efforts by recommending remediation workflows that resolve multiple alerts simultaneously. While using manual correlation or static-based detections can be difficult, consider employing AI to enhance detections and scale resolution efforts. You can also streamline processes with automated playbooks, reduce manual labor, and reduce remediation times.

#4 Secure Sensitive Data

Cloud migration, lift-and-shift operations, and the adoption of microservices technologies for application development are generating more data, in more locations, than ever before. Meanwhile, the increase in data volumes, stringent regulations, and the shift toward AI make it challenging to effectively monitor and protect data in the cloud. Many organizations struggle to understand what sensitive data (such as customer details, health data, or financial information) they actually hold, who can access it, and where it's at risk.

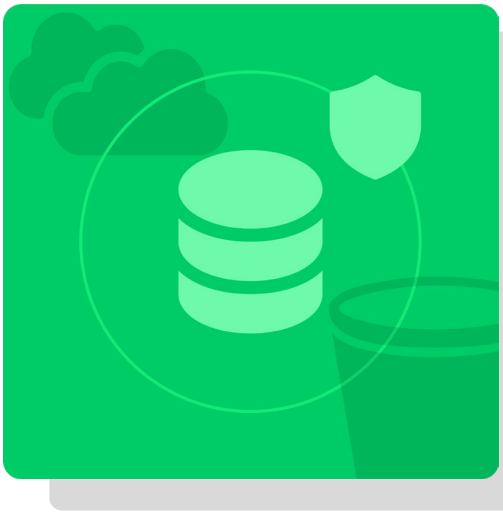

Shadow data is another risk. Every backup and migration increases the risk of moving confidential files to unmonitored storage, or becoming forgotten or abandoned. This can also create more opportunities for threat actors looking to steal data or orchestrate ransomware attacks.

Knowing the data you have and the regulatory frameworks your data adheres to is essential for assessing its risk. From Azure and AWS to Snowflake and BigQuery—as well as across distributed architecture and AI supply chains—your organization needs an up-to-date, accurately classified inventory of sensitive data.

Securing sensitive data in modern cloud deployments requires the use of automated, cloud-native DSPM tools to discover, classify, govern, and mitigate data risks while prioritizing issues by sensitivity and business context. They assess, monitor, and minimize risks related to data that resides in cloud datastores, particularly across multicloud environments. By emphasizing the protection of sensitive information, DSPM examines the context and content of the data and prioritizes PII, medical records, and other critical information.

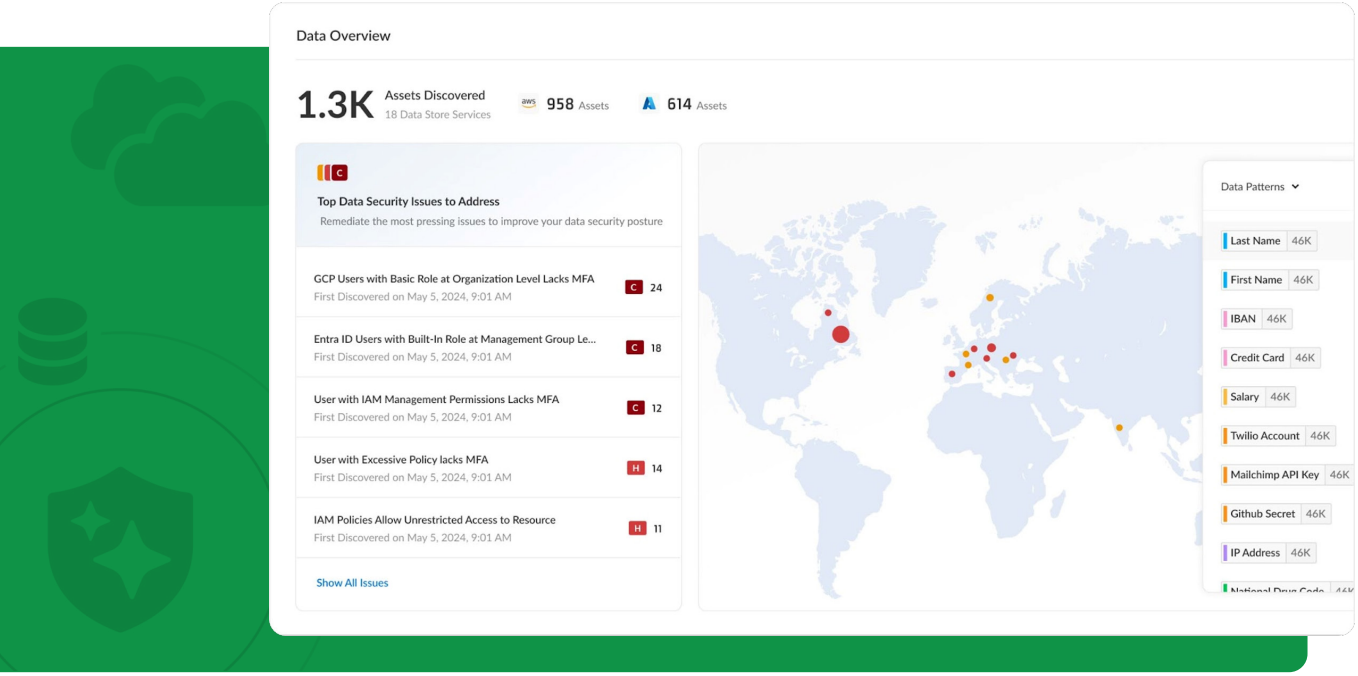

#5 Protect AI Models

AI and LLMs are central to the strategic focus of modern organizations. In recent years, a variety of new AI and LLM tools—including managed model APIs, open-source foundation models, and an expanding ecosystem of operational tooling, plugins, and applications—have emerged. However, while organizations enjoy unprecedented opportunities to build and deploy AI-powered applications rapidly, there's no shortage of accompanying security challenges.

The proliferation of managed, semimanaged, and unmanaged AI models can make oversight difficult. The inner workings of large AI models are also often opaque, even to their creators, making it difficult to anticipate potential security and compliance issues. Models may exhibit unexpected behaviors or vulnerabilities that aren't easily detectable through traditional testing methods.

Existing security measures such as firewalls and posture analysis tools (e.g., CSPM and DSPM) don't address AI-specific attacks like data poisoning, model inversion, and adversarial attempts.

Security teams need visibility into AI deployment to monitor usage effectively and prevent downstream risk—especially since the EU AI Act imposes new requirements around data privacy, algorithmic bias, and explainable AI. It also raises the stakes for noncompliance, with penalties nearly double those of GDPR. Similar legislation is expected in the US and elsewhere.

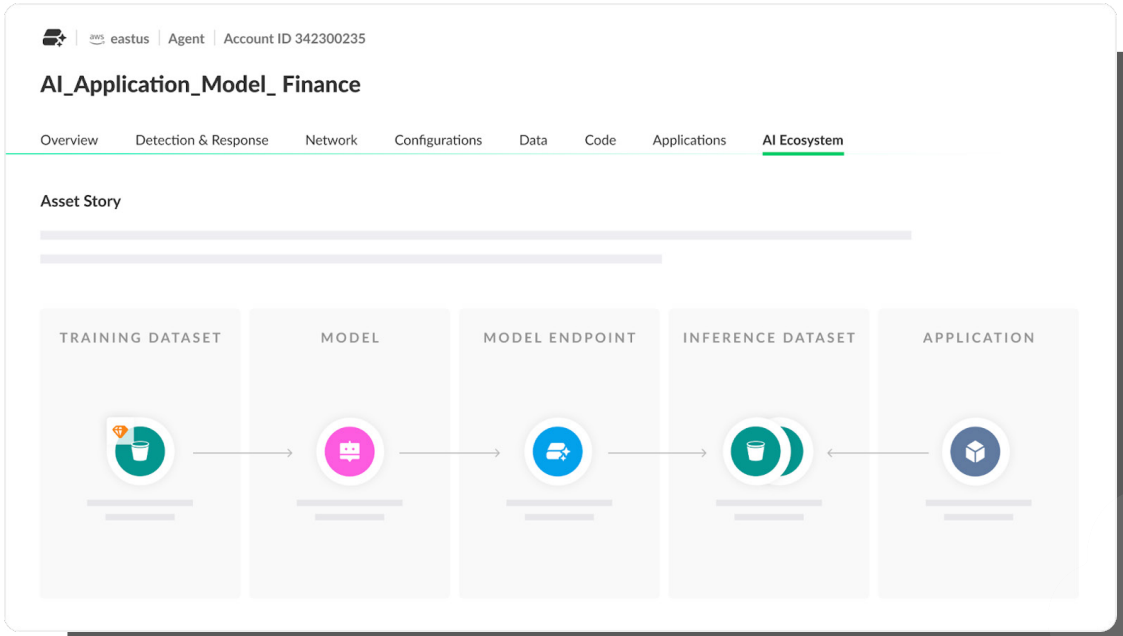

Securing Your Cloud Transformation

Palo Alto Networks Cortex® Cloud secures multicloud environments from code to cloud to SOC. To address the key requirements highlighted in this guide, the solution enables your organization to:

- **Unify Data**
- **Secure Sensitive Data**
- **Detect and Prioritize Risks**
- **Protect AI Models**
- **Remediate Risks**

See Cortex Cloud in Action →

About Cortex Cloud

Cortex® Cloud unites the world's leading CNAPP with the No.1 SecOps platform to deliver real-time security from code to cloud to SOC. Break down silos and gain complete risk context to operate with confidence and unmatched efficiency. Together, unified data, AI, and automation power an adaptive defense that stops threats at their source. Investigate and respond in minutes, and remediate in seconds, with 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

cortex_eb_5-requirements-to-protect-ai-and-cloud-deployments_041625