DUMPSARENA

PSE Palo Alto Networks System Engineer Professional - Prisma Cloud

Palo Alto Networks PSE-PrismaCloud

Version Demo

Total Demo Questions: 14

Total Premium Questions: 149

Buy Premium PDF

https://dumpsarena.co

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

QUESTION NO: 1

Which Prisma Cloud component is installed on a host, container, or Kubernetes node to enforce runtime protection policies?

A. Defender

B. Cloud Account

C. Resource List

D. Intelligence Stream

ANSWER: A

Explanation:

Defender is the Prisma Cloud Compute component that is deployed close to protected workloads. Defenders collect runtime information and enforce the runtime policies configured through the Prisma Cloud Compute Console. Depending on the deployment model, a Defender can protect hosts, containers, Kubernetes workloads, serverless functions, or other supported workload types.

The Console centrally manages policies, visibility, and reporting, while the Defender provides the enforcement point on or alongside the protected workload. Palo Alto Networks describes these architecture components in the Prisma Cloud Compute architecture documentation.

QUESTION NO: 2

Which RQL string searches for all EBS volumes that do not have a "DataClassification" tag?

A. config where api.name = 'aws-ec2-describe-volumes, AND json.rule = tags[*]key contains DataClassification

B. config where api.name = ,aws-ec2-describe-volumes' AND json.rule = tags[*]key != DataClassification

C. config where api.name = ,aws-ec2-describe-volumes' AND json.rule = tags[*].key exists

D. config where api.name = 'aws-ec2-describe-volumes' AND json.rule = tags[*].key = 1

E. config where api.name = 'aws-ec2-describe-volumes' AND json.rule = tags[*].key does not contain 'DataClassification'

ANSWER: E

Explanation:

The query `config where api.name = 'aws-ec2-describe-volumes' AND json.rule = tags[*].key does not contain 'DataClassification'` is correct because it targets the AWS EC2 API data source that returns EBS volume configurations and evaluates the keys in the volume's tag collection. In Prisma Cloud RQL, `tags[*].key` projects the key property from every tag attached to each returned resource. Applying `does not contain 'DataClassification'` identifies volumes for which that collection of tag keys lacks the required tag name. This is the appropriate form for an absence check: it evaluates the complete set of tag keys rather than merely testing whether a tag field exists. The result is a list of EBS volumes that do not carry a tag whose key is exactly `DataClassification`, which can then be used in a config policy to monitor tagging compliance. Palo Alto Networks documents the `config` query structure, API names, JSON rules, and tag-array filtering syntax in the Prisma Cloud RQL Reference. For the AWS resource data model used by Prisma Cloud, see the Config Query Reference.

QUESTION NO: 3

What are two examples of Amazon Web Services logging services? (Choose two.)

A. CloudLog

B. CloudEvent

C. CloudWatch

D. CloudTrail

ANSWER: C D

Explanation:

CloudWatch and **CloudTrail** are AWS services used to collect, retain, monitor, and analyze operational and security-relevant log data. Amazon CloudWatch provides monitoring and observability capabilities, including CloudWatch Logs, which centralizes log events from resources such as Amazon EC2 instances, AWS Lambda functions, containers, and applications. Teams can query logs, configure metric filters, create alarms, and investigate application or infrastructure behavior from this data.

AWS CloudTrail records activity in an AWS environment, particularly API calls and console actions. Its event history and trails provide an auditable record of who performed an action, when it occurred, the source IP address, and the affected AWS resources. This makes CloudTrail especially important for governance, security investigations, compliance evidence, and detecting unexpected configuration changes. In Prisma Cloud deployments, ingesting and analyzing CloudTrail data helps provide visibility into cloud account activity and supports identification of potentially risky actions. Together, CloudWatch and CloudTrail provide complementary logging coverage: CloudWatch focuses on workload and operational telemetry, while CloudTrail focuses on AWS account and control-plane activity. See the Amazon CloudWatch documentation and the AWS CloudTrail User Guide.

QUESTION NO: 4

Which option is true about VM-Series NGFW templates available from the Palo Alto

Networks GitHub repository?

A. Palo Alto Networks provides full support if a valid support license is in place.

B. Support for the templates is available through Professional Services from Palo Alto Networks.

C. Unless otherwise noted, these templates are released under an as-is. best effort support policy.

D. The author of the template provides full support as long as the PAN-OS version specific to the template is supported.

ANSWER: C

Explanation:

Unless otherwise noted, these templates are released under an as-is, best-effort support policy. Palo Alto Networks publishes VM-Series deployment templates and automation examples in its GitHub organizations to help customers deploy and integrate firewalls in cloud environments. These repository artifacts are intended as reference implementations and community resources rather than fully supported product features. Consequently, users should validate a template in their own environment, adapt configuration values to their architecture and security requirements, and perform appropriate testing before using it in production.

The applicable GitHub repository documentation commonly states that its content is provided “as is,” with best-effort support. This designation is important because a valid support subscription for the VM-Series firewall does not automatically make every third-party-style deployment template or repository script subject to the normal TAC support lifecycle. The source repository, its README files, and any template-specific notices should be checked for exceptions or additional support terms. For Palo Alto Networks’ GitHub presence and VM-Series cloud deployment resources, see Palo Alto Networks on GitHub and VM-Series Documentation.

QUESTION NO: 5

Amazon Web Services WAF can be enabled on which two resources?(Choose two.)

A. Amazon CloudFront distribution

B. AWS NAT Gateway

C. AWS ALB

D. AWS NLB

ANSWER: A C

Explanation:

AWS WAF can be associated with an Amazon CloudFront distribution and with an Application Load Balancer. CloudFront is AWS's content delivery network service, and AWS WAF evaluates web requests at the edge before the requests are served from, or forwarded to, an origin. Associating a web ACL with CloudFront provides centralized protection for internet-facing content and APIs delivered through a distribution, including controls such as managed rule groups, IP reputation filtering, rate-based rules, and custom request matching.

An Application Load Balancer is also a supported regional AWS WAF resource. When a web ACL is associated with an ALB, AWS WAF inspects HTTP and HTTPS requests received by the load balancer and applies the configured rule actions before traffic is routed to targets. This makes ALB integration appropriate for protecting web applications hosted behind the load balancer. AWS documents the supported AWS WAF resource types and the process for associating a web ACL with CloudFront distributions or regional resources such as Application Load Balancers in the AWS WAF Developer Guide. See also the AWS WAF overview for its web-application protection model.

QUESTION NO: 6

Which Prisma Cloud Compute capability can block a Kubernetes workload from being deployed when it violates an admission policy?

A. Admission Control

B. Cloud Discovery

C. Host Inventory

D. Compliance Report

ANSWER: A

Explanation:

Admission Control evaluates Kubernetes workload definitions during the admission process before the workload is deployed. It can enforce policies for images, workload configuration, and other deployment criteria, allowing an organization to prevent noncompliant or risky workloads from entering a Kubernetes cluster.

Admission Control provides a preventative control in the deployment lifecycle. Runtime protection is applied after workloads are running, while vulnerability scanning identifies issues in images and hosts. Palo Alto Networks documents Kubernetes admission protection in the Kubernetes Admission Control documentation.

QUESTION NO: 7

Which statement is specific for Prisma Cloud when integrating into cloud environments?

A. An AutoFocus license is included in Prisma Cloud.

B. For multi-cloud environment licenses are required for the number of Prisma Cloud instances.

C. Can be natively integrated into Prisma Access.

D. No agents or proxies are required.

ANSWER: D

Explanation:

No agents or proxies are required. Prisma Cloud can connect to supported cloud environments through agentless onboarding, using cloud-provider APIs and read-only identity permissions to discover accounts, ingest configuration and audit data, and continuously evaluate cloud resources against security policies. This approach enables centralized visibility and posture management without placing a network proxy in the customer traffic path or requiring software installation on the cloud resources being assessed. During onboarding, the organization configures the required cloud account access and then Prisma Cloud retrieves the relevant metadata and telemetry from the provider integration. This is particularly useful for broad, scalable coverage across cloud accounts because security posture can be assessed as soon as the account is connected. Palo Alto Networks documents agentless onboarding as the method used to connect cloud accounts and explains the permissions and integration workflow for supported providers. For workload-specific protection, Prisma Cloud can additionally use Defenders where deeper runtime protection is needed; however, agentless cloud-environment integration itself does not require agents or proxies. See the Prisma Cloud cloud account onboarding documentation and the Prisma Cloud Defender documentation.

QUESTION NO: 8

Which two statements are true about CloudFormation? (Choose two.)

A. CloudFormation is a procedural configuration management tool.

B. CloudFormation templates can be used on both Amazon Web Services and Microsoft Azure

C. CloudFormation templates can be written in JSON or YAML

D. CloudFormation is a declarative orchestration tool.

ANSWER: C D

Explanation:

CloudFormation templates can be written in JSON or YAML because AWS CloudFormation uses these two supported template formats to define the desired state of AWS resources. A template describes resources, their properties, dependencies, parameters, mappings, conditions, and outputs in a structured, machine-readable document. YAML is often preferred for readability and supports CloudFormation-specific short-form intrinsic-function syntax, while JSON remains fully supported for teams or toolchains that standardize on JSON.

CloudFormation is a declarative orchestration tool because an author specifies the intended infrastructure state rather than supplying an ordered imperative procedure for creating every resource. CloudFormation evaluates the template, determines resource dependencies, and provisions, updates, or deletes the stack resources in the necessary order. This model also supports repeatable infrastructure deployments, change sets for reviewing proposed stack updates, and drift detection to identify changes made outside the declared template. These characteristics make CloudFormation suitable for infrastructure as code and consistent lifecycle management of AWS environments. AWS documents the supported template formats and template structure in its CloudFormation template formats documentation, and explains stack creation and management in the AWS CloudFormation User Guide.

QUESTION NO: 9

A client has a sensitive internet-facing application server in Microsoft Azure and is concerned about resource exhaustion because of distributed denial-of-service attacks What can be configured on the VM-Series firewall to specifically protect this server against this type of attack?

A. Custom threat signature

B. Zone Protection Profile

C. QoS Profile to limit incoming requests

D. DoS Protection Profile with specific session counts

ANSWER: D

Explanation:

DoS Protection Profile with specific session counts is correct because it provides targeted, policy-based protection for an individual internet-facing server rather than applying a broad control to an entire zone. On the VM-Series firewall, a DoS Protection profile can be attached to a DoS Protection policy that matches traffic destined for the Azure application server. The profile establishes thresholds for concurrent sessions and/or new sessions per second, enabling the firewall to detect when traffic exceeds the configured acceptable rate and take the configured protective action. This directly addresses resource-exhaustion attacks by limiting excessive connection creation or session consumption before the protected server's resources are overwhelmed. The policy can be configured as classified protection, which allows thresholds to be enforced using traffic attributes such as source and destination IP address, and is especially useful when protection must be focused on a particular critical server. Palo Alto Networks documents DoS Protection as a mechanism for defending resources from floods of sessions or connection attempts and distinguishes it from zone-wide protection controls. See the DoS Protection and Zone Protection documentation and the DoS Protection administration guidance.

QUESTION NO: 10

When an on-premises NGFW (customer gateway) is used to connect to the Virtual

Gateway, which two IKE profiles cannot be used? (Choose two.)

A. Group2 / SHA-1 / AES-128-CBC / IKE-V1

B. Group2 / SHA-1 / AES-128-GCM / IKE-V1

C. Group14 / SHA-256 / AES-256-GCM / IKE-V1

D. Group2 / SHA-1 / AES-128-CBC

E. Group14 / SHA-256 / AES-256-CBC / IKE-V1

ANSWER: C E

Explanation:

Group14 / SHA-256 / AES-256-GCM / IKE-V1 and Group14 / SHA-256 / AES-256-CBC / IKE-V1 cannot be selected for this customer-gateway-to-Virtual-Gateway connection. The relevant compatibility requirement is the complete IKE profile, not merely whether an individual cipher, hash, or Diffie-Hellman group is generally available on a Palo Alto Networks firewall. For IKEv1 interoperability with the Virtual Gateway, the Group 14, SHA-256, and AES-256 profile combination is outside the supported profile set. This remains true whether AES-256 is expressed with GCM or CBC, because both proposed profiles use the unsupported IKEv1 Group 14/SHA-256/AES-256 suite. When configuring the on-premises NGFW, the IKE crypto profile and the peer configuration must use a Virtual-Gateway-supported IKEv1 proposal so that phase 1 negotiation can establish successfully. Palo Alto Networks documents the role of IKE crypto profiles and matching IKE negotiation parameters in its IKE Crypto Profiles documentation; Prisma Access remote-network deployment guidance is available in the Deploy Remote Networks documentation.

QUESTION NO: 11 - (SIMULATION)

A customer has deployed a VM-Series NGFW on Amazon Web Services using a PAYG license. What is the sequence required by the customer to switch to a BYOL license?

ANSWER: Check the explanation for the answer

Explanation:

Required sequence:

1. **Back up** the configuration from the existing PAYG VM-Series firewall.
2. **Deploy** a new VM-Series firewall using the AWS BYOL image.

3. **Register** the new BYOL firewall/authentication code with the Palo Alto Networks support portal.
4. **Activate** the BYOL licenses on the new firewall.
5. **Load** (import and commit) the saved PAYG configuration on the BYOL firewall.

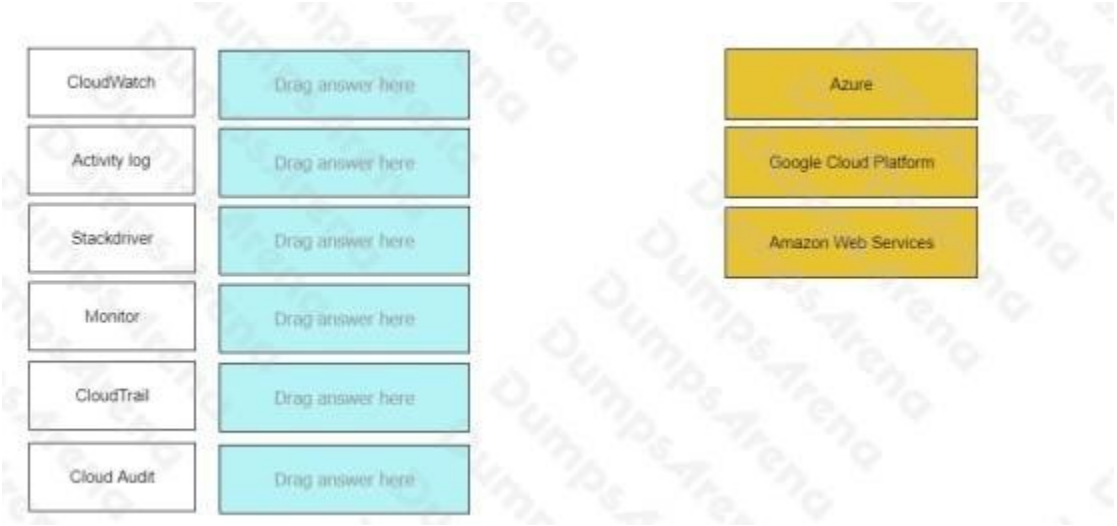
In short: `Backup → Deploy → Register → Activate → Load`.

A PAYG marketplace instance cannot simply be converted in place; the customer deploys a BYOL instance, licenses it, and then restores the configuration.

QUESTION NO: 12 - (DRAG DROP)

DRAG DROP

Match the logging service with its cloud provider.



ANSWER:



Explanation:

The completed mapping is correct because each listed item is the native logging, auditing, or monitoring service associated with its respective cloud platform. Amazon Web Services provides Amazon CloudWatch for collecting, monitoring, and analyzing operational metrics, logs, and events from AWS resources and applications. AWS also provides AWS CloudTrail to record account activity and API actions for governance, compliance, operational auditing, and risk auditing. Palo Alto Networks integrations commonly use these AWS log sources to obtain visibility into cloud activity. AWS documentation describes these services at Amazon CloudWatch and AWS CloudTrail.

In Microsoft Azure, the Activity Log is the platform-level log that records subscription events, while Azure Monitor is Microsoft's monitoring platform for collecting and analyzing telemetry across Azure resources and workloads. These two services therefore belong to Azure. Microsoft documents the Activity Log at Azure Activity Log and the monitoring platform at Azure Monitor overview.

For Google Cloud Platform, Stackdriver is the historical Google Cloud name for the monitoring and logging suite, now represented by Google Cloud Observability services. Cloud Audit Logs records administrative activity, data access, system events, and policy-denied events in Google Cloud. Consequently, Stackdriver and Cloud Audit correctly map to Google Cloud Platform. Google Cloud provides current documentation for Cloud Monitoring and Cloud Audit Logs. The image's final placements consistently reflect these provider-native service associations.

QUESTION NO: 13

Which three features are not supported by VM-Series NGFWs on Azure Stack? (Choose three.)

A. Azure Application Insight

B. Resource Group

C. Azure Security Center

D. Bootstrapping

E. ARM Template

ANSWER: A C D

Explanation:

Azure Application Insight, Azure Security Center, and Bootstrapping are not supported for VM-Series firewalls deployed on Azure Stack. Azure Stack is an extension of Azure designed for on-premises or disconnected/hybrid environments, but its available platform services and VM-Series integration capabilities do not fully match those of public Azure. Consequently, the VM-Series Azure Stack deployment model does not integrate with Azure Application Insight telemetry or Azure Security Center capabilities. It also does not support VM-Series bootstrapping, which on supported cloud platforms automates initial firewall configuration through cloud-init data and associated storage resources. Administrators deploying VM-Series on Azure Stack must instead use the supported Azure Stack deployment and configuration workflow, including the applicable template-based deployment approach and post-deployment firewall configuration. Palo Alto Networks documents these Azure Stack-specific feature limitations in its VM-Series deployment guidance. See the VM-Series on Azure Stack deployment documentation and the VM-Series documentation portal for platform requirements and supported deployment capabilities.

QUESTION NO: 14

Which two resources provide operational insight within the Prisma Cloud Asset Inventory? (Choose two.)

A. Cortex Data Lake

B. Cloud Storage buckets

C. Prisma Access Gateways

D. Compute Engine instance

ANSWER: B D

Explanation:

Cloud Storage buckets and Compute Engine instance are resources that Prisma Cloud can discover, inventory, and contextualize within its cloud-asset visibility capabilities. Asset Inventory provides a centralized view of resources deployed across onboarded cloud accounts, enabling security and operations teams to identify what exists, determine where it is deployed, associate it with its cloud account and project, and investigate its configuration and security posture. In a Google

Cloud environment, Cloud Storage buckets are storage assets whose inventory information is useful for operational visibility and data-exposure investigations. Compute Engine instances are workload assets whose inventory records support visibility into the compute footprint and can be correlated with configuration, network, identity, and risk information. This resource-centric inventory is foundational to Prisma Cloud's ability to help teams search and investigate cloud environments at scale. Palo Alto Networks describes Asset Inventory as the location for viewing cloud resources across connected environments and using resource details for investigation. See the Prisma Cloud Asset Inventory documentation and the Asset Inventory overview.