

# Data Center PSE Professional Practice Test (Sample)

## Study Guide



**Everything you need from our exam experts!**

**Copyright © 2025 by Examzify - A Kaluba Technologies Inc. product.**

**ALL RIGHTS RESERVED.**

**No part of this book may be reproduced or transferred in any form or by any means, graphic, electronic, or mechanical, including photocopying, recording, web distribution, taping, or by any information storage retrieval system, without the written permission of the author.**

**Notice: Examzify makes every reasonable effort to obtain from reliable sources accurate, complete, and timely information about this product.**

**SAMPLE**

## **Questions**

- 1. How can one explain to a customer that their existing firewall may be inadequate?**
  - A. Explain the advantages of App-ID over port filtering.**
  - B. Explain the advantages of App-ID over intrusion prevention.**
  - C. Offer to install a NGFW in TAP mode temporarily.**
  - D. Offer to install a NGFW in Layer 3 mode temporarily.**
- 2. Does traditional high availability incur more costs than resilience on the public cloud, and if so, why?**
  - A. No, the costs are about the same**
  - B. Yes, because in traditional HA you must have a passive VM running at all times**
  - C. Yes, because traditional high availability requires a special license.**
  - D. No, because at low utilization rates, traditional high availability actually is less expensive.**
- 3. Which feature of OpenStack influences service scaling based on traffic metrics?**
  - A. Celiometer**
  - B. Heat Templates**
  - C. Nova Scheduler**
  - D. Neutron Networking**
- 4. Which HEAT template file defines the environment for the VM-Series firewall?**
  - A. pan\_basic\_gw\_env.yaml**
  - B. init-cfg.txt**
  - C. bootstrap.xml**
  - D. pan\_basic\_gw.yaml**
- 5. What do in-band tests measure in the context of firewalls?**
  - A. Throughput, threats, and compatibility**
  - B. Only performance limits**
  - C. Configuration steps**
  - D. Only user access**

- 6. What is the limit enforced for the number of IP addresses per list with EDLs?**
- A. No limits are enforced**
  - B. One limit per list**
  - C. Ten limits per list**
  - D. Twenty limits per list**
- 7. Where do you create Dynamic Address Groups when deploying the Palo Alto Networks NGFW on NSX?**
- A. NSX Manager**
  - B. Panorama**
  - C. Palo Alto Networks NGFW**
  - D. vCenter**
- 8. Can Prisma Cloud scan container images in both public and private repositories?**
- A. No, it is limited to public repositories**
  - B. Yes, but only on public registries**
  - C. Yes, on both public and private repositories**
  - D. No, it only scans local files**
- 9. What governs the interaction of Endpoint Groups in Cisco ACI?**
- A. Policies**
  - B. Contracts**
  - C. Rules**
  - D. Conditions**
- 10. How many nodes can each CN-MGMT service/pair secure?**
- A. 15**
  - B. 20**
  - C. 30**
  - D. 35**

## **Answers**

SAMPLE

1. C
2. B
3. A
4. A
5. A
6. A
7. B
8. C
9. B
10. C

SAMPLE

## **Explanations**

SAMPLE



1. How can one explain to a customer that their existing firewall may be inadequate?
  - A. Explain the advantages of App-ID over port filtering.
  - B. Explain the advantages of App-ID over intrusion prevention.
  - C. Offer to install a NGFW in TAP mode temporarily.**
  - D. Offer to install a NGFW in Layer 3 mode temporarily.

When discussing the inadequacy of an existing firewall, the most effective approach is to demonstrate the capabilities of a next-generation firewall (NGFW) in a practical and low-risk manner. Offering to install the NGFW in TAP mode temporarily allows the customer to see the benefits and enhanced capabilities without making a permanent change to their existing infrastructure. TAP mode enables the NGFW to monitor traffic passively. This means it can analyze and provide insights into the current network traffic, security events, and potential vulnerabilities without interfering with the traffic flow. It helps the customer understand how their existing firewall is functioning and highlights any shortcomings in real-time. By doing so, they can observe the differences and assess how the NGFW would perform in protecting their environment, thereby leading to informed decision-making. This method is particularly compelling because it provides the customer with concrete data and observations that can reinforce the argument for an upgrade or change, rather than just theoretical explanations of features or benefits. The hands-on evaluation experience can be much more persuasive than a purely conversational approach, as customers often respond more positively to direct evidence of value.

2. Does traditional high availability incur more costs than resilience on the public cloud, and if so, why?
  - A. No, the costs are about the same
  - B. Yes, because in traditional HA you must have a passive VM running at all times**
  - C. Yes, because traditional high availability requires a special license.
  - D. No, because at low utilization rates, traditional high availability actually is less expensive.

The assertion that traditional high availability incurs more costs than resilience in the public cloud can be supported by the fact that traditional high availability (HA) usually involves maintaining a redundant infrastructure. In this approach, an organization must have a passive virtual machine (VM) that remains on standby to take over in case the primary VM fails. This means that resources are effectively being wasted during normal operations since the standby system is not actively utilized but still generates costs. In the cloud environment, resilience can be achieved through various design patterns that do not strictly rely on continuous active-passive configurations. Instead, resources can be scaled dynamically based on demand and only provisioned as necessary, leading to potential cost savings. The flexibility of the cloud allows businesses to implement strategies such as failover systems that do not necessitate a constant standby infrastructure, thus reducing ongoing expenses. Understanding the distinction between these models highlights why traditional high availability can lead to higher costs, given that it requires constant resource allocation for the passive instance, whereas resilience strategies leverage the on-demand nature of cloud resources for more efficient and cost-effective operations.

### 3. Which feature of OpenStack influences service scaling based on traffic metrics?

- A. Celiometer**
- B. Heat Templates**
- C. Nova Scheduler**
- D. Neutron Networking**

Celiometer is the component of OpenStack that provides telemetry services, which collect and store measurements of various resources within the OpenStack environment. This feature is integral for influencing service scaling based on traffic metrics because it allows for the monitoring of resource usage, events, and performance data. By gathering these metrics, administrators can make informed decisions about when to scale services up or down to meet varying demand. Effective scaling relies on accurate data regarding the current load and resource usage; thus, Celiometer supports automated scaling mechanisms by providing real-time insights into how services are performing. This telemetry data can trigger scaling activities, ensuring that resources are allocated efficiently in response to traffic spikes or drops. In contrast, Heat Templates are used for orchestration and defining resource stacks but do not directly involve traffic metrics for scaling purposes. Nova Scheduler's primary function is to manage the allocation of compute resources rather than monitoring metrics, and Neutron Networking focuses on managing networking within OpenStack, not directly influencing scaling based on service demand. This makes Celiometer the key feature for scaling services in response to observed traffic metrics.

### 4. Which HEAT template file defines the environment for the VM-Series firewall?

- A. pan\_basic\_gw\_env.yaml**
- B. init-cfg.txt**
- C. bootstrap.xml**
- D. pan\_basic\_gw.yaml**

The correct choice for defining the environment for the VM-Series firewall is the `pan_basic_gw_env.yaml` file. This file specifically contains the configurations and settings necessary for establishing the basic environment in which the VM-Series firewall operates. It outlines important parameters that guide how the firewall interacts with the surrounding infrastructure, such as networking settings and resource allocations. This file is key for automating the deployment of the VM-Series firewall in a cloud or virtualized environment, ensuring that it is correctly configured to meet the organization's security needs. The structured format of YAML allows for clarity and ease of editing, making it suitable for infrastructure-as-code practices. In contrast, other options serve different purposes. For instance, `init-cfg.txt` contains initial configuration commands for the firewall but does not define the environment itself. `Bootstrap.xml` is typically used for the initial bootstrapping process but is not the primary configuration file for setting the environment. Lastly, `pan_basic_gw.yaml` may contain configurations for the gateway firewall but does not specifically outline the environmental parameters as the correct choice does. Thus, the `pan_basic_gw_env.yaml` is essential for delineating the variables that affect the deployment and operation of the firewall in its intended environment.

**5. What do in-band tests measure in the context of firewalls?**

**A. Throughput, threats, and compatibility**

**B. Only performance limits**

**C. Configuration steps**

**D. Only user access**

In the context of firewalls, in-band tests are designed to assess how well the firewall performs under various conditions and configurations. The measurement of throughput indicates how much data can be processed by the firewall within a given period, which reflects its efficiency and capability to handle network traffic. Evaluating threats is crucial to ensure that the firewall can effectively detect and mitigate various types of cyber threats, such as intrusions or malicious activities, while ensuring safe communication across the network. Compatibility testing ensures that the firewall can work with existing network infrastructures and other security tools without causing conflicts or impairing performance. Overall, the comprehensive nature of in-band tests allows for a holistic assessment of firewall performance, security efficacy, and operational compatibility, making this option the most accurate representation of what in-band tests measure.

**6. What is the limit enforced for the number of IP addresses per list with EDLs?**

**A. No limits are enforced**

**B. One limit per list**

**C. Ten limits per list**

**D. Twenty limits per list**

The correct response indicates that no limits are enforced on the number of IP addresses that can be included in an External Dynamic List (EDL). EDLs are designed for dynamic management of IP addresses, allowing for flexible and scalable configurations based on evolving network needs or external data sources. This means organizations can utilize EDLs to adapt to various scenarios without being constrained by a fixed number of entries, making them particularly useful for security measures such as threat intelligence sharing or access control based on frequently changing data. Understanding that there are no imposed limits on the number of IP addresses allows users to leverage EDLs more effectively. This flexibility can be crucial for maintaining security in environments where the threat landscape is constantly changing, as it enables quick updates and modifications based on real-time data.

**7. Where do you create Dynamic Address Groups when deploying the Palo Alto Networks NGFW on NSX?**

- A. NSX Manager**
- B. Panorama**
- C. Palo Alto Networks NGFW**
- D. vCenter**

Dynamic Address Groups are a crucial feature used in conjunction with the Palo Alto Networks Next-Generation Firewall (NGFW) when deployed in an NSX environment. They enable more granular and flexible security policies by dynamically including virtual machines based on criteria such as IP address, tags, or other attributes. Creating Dynamic Address Groups typically requires a centralized management approach, which is why Panorama is the correct choice in this context. Panorama serves as a management platform for multiple Palo Alto firewalls and provides visibility, centralized policy management, and monitoring. It facilitates the creation and management of dynamic address groups across the entire infrastructure, ensuring that security policies can dynamically adjust as the underlying virtual machine infrastructure changes. While NSX Manager, vCenter, and the Palo Alto Networks NGFW have their respective roles in the overall management and configuration of the environment, they do not provide the same centralized policy management functionality required for creating and managing dynamic address groups. NSX Manager and vCenter are primarily concerned with network and virtualization management, while the NGFW handles firewall policies primarily. Thus, Panorama is the most suitable and designed tool for this task, allowing better integration and management within the Palo Alto security framework.

**8. Can Prisma Cloud scan container images in both public and private repositories?**

- A. No, it is limited to public repositories**
- B. Yes, but only on public registries**
- C. Yes, on both public and private repositories**
- D. No, it only scans local files**

Prisma Cloud is designed to provide comprehensive security for cloud-native applications, which includes the ability to scan container images. When considering its capabilities, it's important to recognize that it can indeed scan both public and private repositories. This feature allows organizations to ensure that their container images are secure regardless of where they are stored, thus reinforcing the security posture across different environments. By scanning images in both public and private repositories, Prisma Cloud enables users to identify vulnerabilities, compliance issues, and other security risks at any point in their deployment pipeline. This is particularly critical for private repositories where sensitive or proprietary software may be stored, making it essential to maintain security standards to protect against potential threats. Therefore, the correct answer reflects the versatility and robustness of Prisma Cloud's scanning capabilities, emphasizing its effectiveness in a variety of settings, ensuring that security is not compromised whether images are in a public or private storage environment.

## 9. What governs the interaction of Endpoint Groups in Cisco ACI?

A. Policies

**B. Contracts**

C. Rules

D. Conditions

The interaction of Endpoint Groups (EPGs) in Cisco ACI is governed by contracts. In Cisco ACI, a contract defines the rules and policies that specify how communication should occur between different EPGs, facilitating the management of security and connectivity. This mechanism allows you to control which EPGs can communicate with each other and under what conditions, enabling fine-tuned access control and segmentation within the network. Contracts can include filters that define which types of traffic are permitted or denied between the EPGs, further enhancing the security posture of the deployment. By using contracts, administrators can create an intentional interaction model based on application requirements and security policies, ensuring that only authorized traffic flows between endpoints. This focus on contractual agreements aligns with ACI's overarching design philosophy, which emphasizes policy-driven automation and includes capabilities for programmability and orchestration. The distinct approach of using contracts instead of more traditional methods ensures that interactions are intentional and manageable within the context of the overall network architecture.

## 10. How many nodes can each CN-MGMT service/pair secure?

A. 15

B. 20

**C. 30**

D. 35

The correct answer indicates that each CN-MGMT service pair can secure 30 nodes. Understanding the role of the CN-MGMT (Control Node Management) service in a broader context is essential. In data center environments, CN-MGMT services are responsible for managing and securing nodes, which refer to individual computing units or devices within the infrastructure. By specifying a limit of 30 nodes, it establishes a framework for managing resources effectively. This is crucial because it allows for optimal performance without overloading the management systems, ensuring scalability while maintaining operational efficiency. When managing a large number of nodes, having a defined maximum helps in planning capacity and ensuring that each node receives adequate attention and management from the service. The specified limit also ensures that network latencies remain low and that the system can maintain a high level of responsiveness when executing tasks such as monitoring, updates, and troubleshooting across the managed nodes. Understanding this configuration allows data center professionals to make informed decisions about scaling infrastructure and adopting best practices in resource management.