DUMPSBOSS.

CompTIA CloudNetX Exam

CompTIA CNX-001

Version Demo

Total Demo Questions: 10

Total Premium Questions: 105

Buy Premium PDF

https://dumpsboss.co

support@dumpsboss.co

support@dumpsboss.co
dumpsboss.co

Topic Break Down

Topic	No. of Questions
Topic 1, Cloud Network Architecture	32
Topic 2, Cloud Network Security	36
Topic 3, Cloud Network Operations	8
Topic 4, Cloud Network Troubleshooting	23
Topic 5, Cloud Network Automation and Orchestration	6
Total	105

QUESTION NO: 1

An architect needs to deploy a new payroll application on a cloud host. End users' access to the application will be based on the end users' role. In addition, the host must be deployed on the 192.168.77.32/30 subnet. Which of the following Zero Trust elements are being implemented in this design? (Choose two.)

A. Least privilege

B. Device trust

C. Microsegmentation

D. CASB

E. WAF

F. MFA

ANSWER: A C

Explanation:

Least privilege is implemented through role-based access to the payroll application. In a Zero Trust architecture, access decisions are granted at the smallest practical level and are based on the user's required job functions. Assigning permissions according to each end user's role limits each identity to the application capabilities and payroll data necessary for its authorized responsibilities. This reflects the Zero Trust principle of explicitly authorizing access rather than providing broad, standing permissions.

Microsegmentation is implemented by requiring the cloud host to reside in the specified /30 subnet. A /30 IPv4 network provides four addresses, with two generally usable host addresses, creating a very small and controlled network boundary. Placing the payroll workload in this tightly scoped segment supports granular isolation, permits targeted security-policy enforcement, and reduces the pathways available for lateral movement. Microsegmentation applies protection as close as possible to individual workloads and their communications, rather than relying only on a broad perimeter network.

Together, role-scoped authorization and workload-level network isolation apply complementary Zero Trust controls for both identity access and application hosting. NIST describes Zero Trust as protecting resources through explicit, dynamic access decisions and minimizing implicit trust: NIST SP 800-207. CISA also identifies microsegmentation as an approach for reducing lateral movement in Zero Trust environments: CISA Zero Trust Maturity Model.

QUESTION NO: 2

A security architect needs to increase the security controls around computer hardware installations. The requirements are:

Auditable access logs to computer rooms

Alerts for unauthorized access attempts

Remote visibility to the inside of computer rooms

Which of the following controls best meet these requirements? (Choose two.)

A. Video surveillance

B. NFC access cards

C. Motion sensors

D. Locks and keys

E. Security patrols

F. Automated lighting

ANSWER: A B

Explanation:

Video surveillance and **NFC access cards** together provide the required layered physical-security controls for computer rooms. Video surveillance gives authorized administrators remote, real-time visibility into the room and retains recorded footage for incident investigation and audit evidence. A modern camera system can also integrate with video analytics or a physical access-control platform to generate alerts when it detects activity in restricted areas or outside approved time windows.

NFC access cards provide identity-based entry control rather than simply securing a door. When a badge is presented, the access-control system records the credential, reader location, date, time, and the allow-or-deny result. These centrally retained records create auditable access logs. The same system can alert security staff when a credential is rejected, used outside its authorized schedule, or presented repeatedly at a restricted door. Correlating badge events with camera footage also improves the ability to validate who entered a computer room and investigate attempted unauthorized access.

This combination aligns with defense-in-depth practices: electronic access control supplies attributable entry records and access-attempt alerts, while surveillance supplies remote observation and supporting evidence. See CompTIA CloudNetX and NIST SP 800-53 Rev. 5.

QUESTION NO: 3

A network engineer is designing a Layer 2 deployment for a company that occupies severalfloors in an office building. The engineer decides to make each floor its own VLAN but still allow for communication between all user VLANs. The engineer also wants to reduce the time necessary for STP convergence to occur when new switches come online. Which of the following should the engineer enable to accomplish this goal?

A. BPDU Guard

B. Priority

C. Tagging

D. Portfast

ANSWER: D

Explanation:

PortFast is the applicable STP optimization because it causes a configured edge/access port to transition directly to the forwarding state rather than progressing through the normal STP listening and learning behavior. This makes devices connected to those ports available much more quickly after a link becomes active, avoiding the delay that would otherwise affect user connectivity, DHCP requests, and similar startup traffic. In a floor-based VLAN design, PortFast is normally applied to access interfaces serving endpoint devices in each user VLAN; VLAN separation does not prevent its use. Inter-VLAN communication is separately provided by Layer 3 routing, such as SVIs on a multilayer switch or a router, while PortFast addresses the Layer 2 port-forwarding delay. Cisco documents that PortFast-enabled interfaces immediately move to the STP forwarding state when brought up. It should be deployed only where the interface is expected to connect to an end device, and operational safeguards such as BPDU Guard are commonly paired with it to protect the topology if an unexpected bridge is connected. See Cisco's Understanding STP PortFast, BPDU Guard, and Related Features and the Cisco IOS XE Spanning Tree Configuration Guide.

QUESTION NO: 4

A network architect is designing a solution to secure the organization's applications based on the security policy. The requirements are:

Users must authenticate using one set of credentials.

External users must be located in authorized sites.

Session timeouts must be enforced.

Network access requirements should be changed as needed.

Which of the following best meet these requirements? (Choose two.)

A. Role-based access

B. Single sign-on

C. Static IP allocation

D. Multifactor authentication

E. Conditional access policy

F. Risk-based authentication

ANSWER: B E

Explanation:

Single sign-on meets the requirement for users to authenticate with one set of credentials. An SSO implementation centralizes authentication through an identity provider, allowing a user to establish a session once and then access multiple approved applications without separately authenticating to each one. This supports consistent identity management and reduces the number of credentials that users must manage. CompTIA's CloudNetX objectives include implementing identity and access management capabilities such as federated identity and SSO. See the CompTIA exam objectives resource.

Conditional access policy meets the location, session, and adaptable-access requirements. Conditional access evaluates contextual signals at sign-in or during access, including a user's geographic or named-network location, application, device state, and risk. Policies can limit external access to approved locations, apply sign-in frequency or session controls to enforce timeouts, and be updated as organizational access requirements change. For example, Microsoft Entra Conditional Access supports location-based conditions and session controls that can require periodic reauthentication. See Microsoft Learn: Conditional Access conditions.

QUESTION NO: 5

A network load balancer is not correctly validating a client TLS certificate. The network architect needs to validate the certificate installed on the load balancer before progressing. Which of the following commands should the architect use to confirm whether the private key and certificate match?

A. openssl-list -noout -modulus -in cert.crt | openssl md5openssl rsa -noout -modulus -in privkey.txt | openssl md5

B. openssl req -in certificate.csr -verifyopenssl-verify -noout -modulus -in privkey.txt | openssl md5

C. openssl-rsa -noout -modulus -in cert.crt | openssl md5openssl-verify -noout -modulus -in privkey.txt | openssl md5

D. openssl x509 -noout -modulus -in cert.crt | openssl md5openssl rsa -noout -modulus -in privkey.txt | openssl md5

ANSWER: D

Explanation:

`openssl x509 -noout -modulus -in cert.crt | openssl md5` followed by `openssl rsa -noout -modulus -in privkey.txt | openssl md5` is correct because it obtains a comparable identifier from the RSA public certificate and the corresponding RSA private key. The `openssl x509` command reads the X.509 certificate, and its `-modulus` option outputs the RSA modulus embedded in the certificate's public key. The `openssl rsa` command reads the RSA private-key file and outputs that key's modulus. Piping each value to a digest command makes the comparison convenient: matching digest output confirms that both files contain the same modulus and therefore are the matching certificate/key pair. This check is useful when troubleshooting TLS configuration on a load balancer because a server or client certificate must be paired with its associated private key to complete TLS authentication successfully. The certificate and key file paths and names may differ in an actual deployment, but the validation method remains the same. OpenSSL documents the certificate command and its modulus output in the openssl-x509 reference, while private-key processing and modulus output are covered in the openssl-rsa reference.

QUESTION NO: 6

A network architect is choosing design options for a new SD-WAN installation that has the following requirements:

All network traffic from the cloud must pass through inspection devices in a dedicated data center.

Ensure redundancy.

Centralize egress traffic.

Which of the following network topologies best meets these requirements?

A. Point-to-point

B. Hub-and-spoke

C. Star

D. Partial mesh

ANSWER: B

Explanation:

Hub-and-spoke is the appropriate SD-WAN topology because it directs traffic from cloud-connected sites or remote spokes to a designated central hub before it reaches external destinations. In this design, the dedicated data center acts as the hub and hosts the required inspection stack, such as next-generation firewalls, intrusion prevention systems, secure web gateways, or other policy-enforcement devices. Routing cloud traffic through that hub ensures inspection is consistently applied and prevents sites from using uncontrolled local internet egress paths.

The model also centralizes egress: outbound traffic exits through the data center's controlled internet connections rather than through separate connections at each cloud or branch location. Redundancy can be built into the design by deploying multiple hubs, redundant inspection devices, diverse WAN transports, and resilient tunnels from spokes to each hub. SD-WAN's centralized orchestration and policy-based path selection can then steer traffic to an available hub while preserving the required service-chaining and inspection policy.

This aligns with CloudNetX objectives covering secure connectivity, traffic flow, high availability, and centralized management in hybrid and multicloud networks. See the CompTIA exam objectives resource and Cisco's overview of SD-WAN architecture and centralized policy.

QUESTION NO: 7 - (SIMULATION)

You are designing a campus network with a three-tier hierarchy and need to ensure secure connectivity between locations and traveling employees.

INSTRUCTIONS

Review the command output by clicking on the server, laptops, and workstations on the network.

Use the drop-down menus to determine the appropriate technology and label for each layer on the diagram. Options may only be used once.

Click on the magnifying glass to make additional configuration changes.

If at any time you would like to bring back the initial state of the simulation, please click the Reset All button.

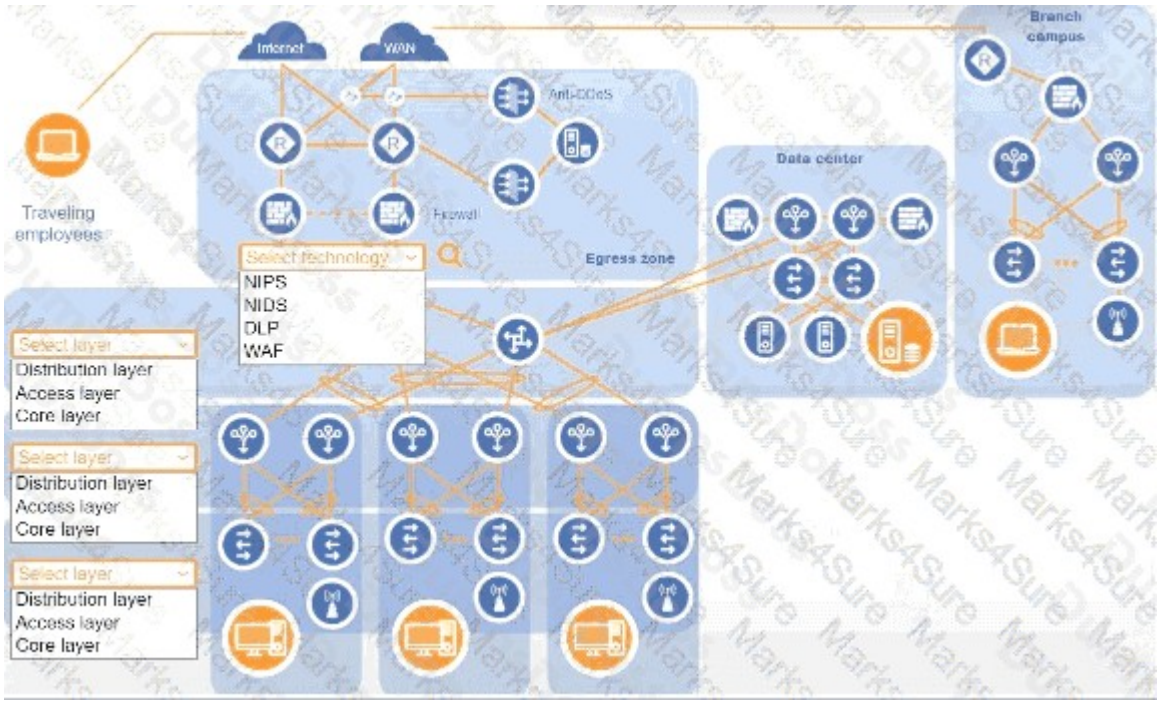

ANSWER: See the explanation for the answer

Explanation:

Select the following values in the simulation:

Security technology: `NIPS`
Top layer dropdown: `Core layer`
Middle layer dropdown: `Distribution layer`
Bottom layer dropdown: `Access layer`

A **NIPS** is appropriately placed with the firewall/perimeter security controls to inspect traffic that has passed the firewall and prevent detected malicious activity. The campus hierarchy proceeds from the backbone/core at the top, through the aggregation/distribution layer, to the access layer where user workstations and other endpoint devices connect.

QUESTION NO: 8

A network architect needs to design a solution to ensure every cloud environment network is built to the same baseline. The solution must meet the following requirements:

Use automated deployment.

Easily update multiple environments.

Share code with a community of practice.

Which of the following are the best solutions? (Choose two.)

A. CI/CD pipelines

B. Public code repository

C. Deployment runbooks

D. Private code repository

E. Automated image deployment

F. Deployment guides

ANSWER: A B

Explanation:

CI/CD pipelines and a **Public code repository** together provide a repeatable, collaborative approach for delivering a standardized cloud-network baseline. A CI/CD pipeline automates the validation and deployment process for infrastructure configuration and network code. When a baseline is defined as version-controlled infrastructure-as-code, the pipeline can apply the same tested configuration consistently to each target cloud environment. Changes to that baseline can be validated and promoted through controlled stages, allowing updates to be delivered efficiently across multiple environments while reducing configuration drift and manual implementation variation.

A Public code repository provides the collaboration and version-control foundation needed to share the baseline templates, modules, pipeline definitions, and documentation with a community of practice. Contributors can review proposed changes, track revisions, reuse approved components, and work from a common source of truth. This combination supports the automation, repeatability, consistency, and knowledge sharing expected in cloud-network operations. CompTIA identifies automation and infrastructure-as-code capabilities as central CloudNetX skills; see the CompTIA CloudNetX certification page. For practical version-controlled automation workflows, see GitHub's CI/CD workflow documentation.

QUESTION NO: 9

A network engineer at an e-commerce organization must improve the following dashboard due to a performance issue on the website:

(Refer to the image: Website performance monitoring dashboard showing metrics like network usage, CPU usage, memory usage, and disk usage over time.)

Which of the following is the most useful information to add to the dashboard for the operations team?

A. 404 errors

B. Concurrent users

C. Number of orders

D. Number of active incidents

ANSWER: B

Explanation:

Concurrent users is the most useful addition because it provides the workload context needed to interpret the existing infrastructure metrics. CPU, memory, disk, and network utilization can show that resources are under pressure, but they do not by themselves indicate whether that pressure corresponds to increased customer demand. Tracking the number of users simultaneously interacting with the e-commerce site lets the operations team correlate traffic volume with changes in response time, errors, and resource consumption. For example, a CPU or network spike that begins as concurrent-user counts rise can indicate a capacity, scaling, application-efficiency, or network-throughput issue. Conversely, poor performance when concurrent use is stable may direct investigation toward an application deployment, dependency, or infrastructure fault rather than demand growth.

This is especially valuable for an e-commerce service because user demand can vary substantially during promotions, product launches, and seasonal events. A concurrent-user metric supports capacity planning and helps teams establish meaningful alert thresholds based on actual service load rather than resource utilization alone. Google's monitoring guidance emphasizes measuring both service behavior and the workload driving it: Google SRE: Monitoring Distributed Systems. Microsoft also describes monitoring application usage and user behavior to understand application load and performance: Azure Monitor Application Insights usage analysis.

QUESTION NO: 10

A company provides an API that runs on the public cloud for its customers. A fixed number of VMs host the APIs. During peak hours, the company notices a spike in usage that results in network communication speeds slowing down for all customers. The management team has decided that access for all customers should be fair and accessible at all times. Which of the following is the most cost-effective way to address this issue?

A. Use an allow list for customers using APIs.

B. Increase the number of VMs running APIs.

C. Enable throttling on APIs.

D. Increase the MTU on the VMs.

ANSWER: C

Explanation:

Enable throttling on APIs is the most cost-effective solution because API throttling establishes enforceable limits on the volume or rate of requests accepted from an individual customer, client, API key, or source during a defined interval. During peak periods, these limits prevent a small number of high-volume consumers from monopolizing the fixed VM-hosted API capacity. The result is more predictable performance and equitable availability across the customer base, directly meeting the requirement for fair access at all times.

Throttling is typically configured at an API gateway, load balancer, web application firewall, or the API-management platform. Common policies include requests per second, requests per minute, concurrent-request limits, quotas, and short burst limits. When a client exceeds its permitted rate, the service can return an HTTP 429 Too Many Requests response, often with a Retry-After header, allowing the client to retry responsibly. This controls demand without requiring the company to continually add compute capacity merely to accommodate transient spikes. API rate-limiting behavior and the 429 response are defined in RFC 6585, Section 4. For a practical cloud implementation example, see AWS API Gateway request throttling documentation.