DUMPSBOSS.

CompTIA Cloud+

CompTIA CV0-004

Version Demo

Total Demo Questions: 15

Total Premium Questions: 258

Buy Premium PDF

https://dumpsboss.co

support@dumpsboss.co

support@dumpsboss.co
dumpsboss.co

QUESTION NO: 1

A security analyst reviews the daily logs and notices the following suspicious activity:

Host	NA/US/John Smith
IP	10.150.71.151
Activity	A powershell process executed compressed, encoded command line content.

The analyst investigates the firewall logs and identities the following:

Operating system	Kali Linux
CPU	x64
Filesystem	ext4
User	John Smith
Category	Compromised - Unauthorized Access
Domain	NA/US
IP	201.101.25.121 (External)
Port	4444
Connection type	Inbound Connection

Which of the following steps should the security analyst take next to resolve this issue? (Select two).

A. Submit an IT support ticket and request Kali Linux be uninstalled from John Smith's computer

B. Block all inbound connections on port 4444 and block the IP address 201.101.25.121.

C. Contact John Smith and request the Ethernet cable attached to the desktop be unplugged

D. Check the running processes to confirm if a backdoor connection has been established.

E. Upgrade the Windows x64 operating system on John Smith's computer to the latest version.

F. Block all outbound connections from the IP address 10.150.71.151.

ANSWER: B D

Explanation:

The PowerShell alert is a big red flag because “compressed/encoded command line content” is a common trick used by attackers to hide what a script is really doing. Pair that with a firewall log showing an inbound connection from an external IP on port 4444 (a port frequently used by reverse shells and tools like Metasploit), and you have a strong indicator of an active compromise.

The next best move is to stop the bleeding and confirm what’s running. Blocking inbound traffic on port 4444 and blocking the specific source IP (201.101.25.121) helps cut off the attacker’s current access path while you investigate. Then, checking the running processes on the affected host helps you verify whether a backdoor, remote shell, or other malicious process is already established and persisting.

Options like “uninstall Kali Linux” or “upgrade Windows” don’t address the immediate threat, and blocking all outbound connections from the internal host is usually too disruptive without confirmation. For general incident response guidance (containment and analysis), see https://www.cisa.gov/resources-tools/resources/incident-handlers-handbook and MITRE technique notes on encoded PowerShell at https://attack.mitre.org/techniques/T1059/001/.

QUESTION NO: 2 - (HOTSPOT)

HOTSPOT

A highly regulated business is required to work remotely, and the risk tolerance is very low. You are tasked with providing an identity solution to the company cloud that includes the following:

secure connectivity that minimizes user login

tracks user activity and monitors for anomalous activity

requires secondary authentication

INSTRUCTIONS

Select controls and servers for the proper control points.

ANSWER:

Explanation:

VPN Client

MFA

SIEM

QUESTION NO: 3

A customer is migrating applications to the cloud and wants to grant authorization based on the classification levels of each system. Which of the following should the customer implement to ensure authorisation to systems is granted when the user and system classification properties match? (Select two).

A. Resource tagging

B. Discretionary access control

C. Multifactor authentication

D. Role-based access control

E. Token-based authentication

F. Bastion host

ANSWER: B D

Explanation:

To grant access based on “classification levels” (think labels like Public, Confidential, Secret), you’re really looking for access control methods that can map user permissions to those labels in a consistent way.

Role-based access control (RBAC) is a solid fit because you can build roles around classification needs (for example, “Confidential-Approved Users”) and then assign users to those roles. Once the role is set, access is automatically enforced across systems that share that classification requirement.

Discretionary access control (DAC) can also work in environments where the data/system owner is allowed to decide who gets access. If the owner is assigning permissions based on classification rules, DAC supports that kind of per-resource decision-making.

Options like MFA and token-based authentication help prove who you are (authentication), but they don’t decide what you’re allowed to access (authorization). Resource tagging is useful for labeling resources, but by itself it doesn’t enforce access unless it’s paired with a policy engine.

References: https://csrc.nist.gov/glossary/term/role_based_access_control and https://csrc.nist.gov/glossary/term/discretionary_access_control

QUESTION NO: 4

A cloud engineer wants to run a script that increases the volume storage size if it is below 100GB.

Which of the following should the engineer run?

A.
```
if [ VOL = describe_volume_size(get_volume(VM)) < 100]
    resize_size(VOL)
else
     echo "$vol is already larger than 100GB"
```

B.
```
if [ VOL = describe_volume_size(get_volume(VM)) + 100]
    resize_size(VOL)
else
     echo "$vol is already larger than 100GB"
```

C.
```
if [ VOL = describe_volume_size(get_volume(VM)) != 100]
    resize_size(VOL)
else
     echo "$vol is already larger than 100GB"
```

D.
```
if [ VOL = describe_volume_size(get_volume(VM)) == 100]
    resize_size(VOL)
else
     echo "$vol is already larger than 100GB"
```

A. Option A

B. Option B

C. Option C

D. Option D

ANSWER: A

Explanation:

Option A is the only one that matches the requirement: “increase the volume size if it is below 100GB.” The key part is the comparison operator `< 100`, which checks whether the current volume size is under the threshold. If it is, the script calls `resize_size(VOL)` to grow it; otherwise it prints a message and does nothing.

The other options don’t line up with the goal. Option B adds 100 to the current size, which would resize even when the volume is already bigger than 100GB (and it’s not even a real conditional check). Option C resizes whenever the size is anything other than exactly 100GB, which is way too broad. Option D only resizes when the size is exactly 100GB, which is the opposite of what you want.

In real cloud scripting (AWS/GCP/Azure), this same idea is usually done by querying the current disk size, comparing it to a target, and then calling the provider’s “modify/resize disk” action only when the current size is smaller. For a similar concept, see https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-modify-volume.html

QUESTION NO: 5

A retail store is rolling out a new point-of-sale solution to several locations where the staff members normally use cash registers. Which of the following activities will best assist staff members with integrating this new solution? (Select two).

A. Documentation

B. Team collaboration

C. Gap analysis

D. Training sessions

E. Touch points

F. Progress monitoring

ANSWER: A D

Explanation:

The two best choices are **documentation** and **training sessions** because they directly help cash-register users learn and use the new POS system day to day. Good documentation (quick-start guides, step-by-step checkout/refund procedures, troubleshooting tips) gives staff something to fall back on when they get stuck during a busy shift.

Training sessions are just as important because they let employees practice real tasks—ringing up items, handling returns, opening/closing a register, and dealing with common errors—before they're live with customers. That hands-on time reduces mistakes and frustration, which is usually what makes or breaks adoption.

Things like gap analysis and progress monitoring are useful for the project team, but they don't directly teach staff how to operate the new system. For more on change enablement and training concepts, see https://www.comptia.org/training/resources and for practical user documentation guidance, see https://www.atlassian.com/software/confluence/guides/productivity/how-to-write-documentation

QUESTION NO: 6

Which of the following cloud deployment models is the best way to replicate a workload nondisruptively between on-premises servers and a public cloud?

A. Public

B. Community

C. Private

D. Hybrid

ANSWER: D

Explanation:

The best fit here is a **hybrid cloud** because it's specifically designed to connect your on-prem environment with a public cloud in a coordinated way. That connection (usually through VPNs, dedicated links, identity integration, and consistent networking) lets you keep systems running while you replicate data and workloads back and forth.

In real life, this is how teams do things like live replication, warm standby, disaster recovery, or gradual migrations without taking an outage. You can sync data from on-prem to the public cloud, test and validate in the cloud, and then cut over when you're ready—without disrupting the original workload.

Public-only, private-only, or community models don't naturally solve the "between on-prem and public cloud" requirement, because they don't describe an integrated environment spanning both. Hybrid is the model that explicitly covers that mixed setup.

References: https://www.nist.gov/publications/nist-definition-cloud-computing and https://learn.microsoft.com/en-us/azure/cloud-adoption-framework/scenarios/hybrid

QUESTION NO: 7

A software engineer is integrating an application lo The cloud that is web socket based. Which of the following applications is the engineer most likely deploying?

A. Image-sharing

B. Data visualization

C. Chat

D. File transfer

ANSWER: C

Explanation:

WebSockets are mainly used when you need real-time, two-way communication between a client and a server over one long-lived connection. Instead of the client constantly "checking in" (polling) for updates, the server can push messages instantly as soon as something happens.

That behavior lines up perfectly with a chat app: users expect messages to appear immediately, typing indicators can update live, and the connection needs to stay open to keep the conversation flowing smoothly. WebSockets are basically built for that kind of always-on, interactive experience.

The other choices don't usually require a persistent, full-duplex connection. Image-sharing and file transfer are more commonly handled with standard HTTP/HTTPS uploads and downloads, and data visualization is often fine with normal API calls unless it's doing live streaming dashboards.

References: https://developer.mozilla.org/en-US/docs/Web/API/WebSockets_API and https://datatracker.ietf.org/doc/html/rfc6455

QUESTION NO: 8

A cloud engineer needs to deploy a new version of a web application to 100 servers. In the past, new version deployments have caused outages. Which of the following deployment types should the cloud engineer implement to prevent the outages from happening this time?

A. Rolling

B. Blue-green

C. Canary

D. Round-robin

ANSWER: C

Explanation:

A canary deployment is the safest fit here because it lets you roll out the new version to a small slice of servers (or users) first, instead of blasting it to all 100 servers at once. If something goes wrong—bad config, hidden bug, performance issue—you'll see it early while most users are still on the stable version.

The big win is controlled risk. You can monitor error rates, latency, and resource usage on the canary group, and if things look bad, you roll back quickly without taking down the whole app. If everything looks good, you gradually expand the rollout until all servers are updated. That "test in production, but safely" approach is exactly what helps avoid the full-scale outages they've seen before.

References: https://martinfowler.com/bliki/CanaryRelease.html and https://kubernetes.io/docs/concepts/cluster-administration/manage-deployment/#canary-deployments

QUESTION NO: 9

Which of the following storage resources provides higher availability and speed for currently used files?

A. Warm/HDD

B. Cold/SSD

C. Hot/SSD

D. Archive/HDD

ANSWER: C

Explanation:

For files that are being used right now (or accessed all the time), you typically want the "hot" tier. Hot storage is built for frequent reads and writes, so it's designed to stay highly available and respond quickly when apps or users request data.

Pairing that hot tier with SSDs makes it even faster. SSDs have much lower latency and higher IOPS than spinning hard drives, so they're a better fit when performance matters—like active project files, databases, or VM disks.

By contrast, warm or cold tiers are meant for less frequently accessed data, and archive is for data you almost never touch. Those tiers can be cheaper, but they usually trade off speed (and sometimes availability options) because the data just isn't needed as often.

References: https://aws.amazon.com/s3/storage-classes/ and https://learn.microsoft.com/en-us/azure/storage/blobs/access-tiers-overview

QUESTION NO: 10 - (SIMULATION)

A company has decided to scale its e-commerce application from its corporate datacenter to a commercial cloud provider to meet an anticipated increase in demand during an upcoming holiday.

The majority of the application load takes place on the application server under normal conditions. For this reason, the company decides to deploy additional application servers into a commercial cloud provider using the on-premises orchestration engine that installs and configures common software and network configurations.

The remote computing environment is connected to the on-premises datacenter via a site-to-site IPSec tunnel. The external DNS provider has been configured to use weighted round-robin routing to load balance connections from the Internet.

During testing, the company discovers that only 20% of connections completed successfully.

INSTRUCTIONS

Review the network architecture and supporting documents and fulfill these requirements:

Part 1:

Analyze the configuration of the following components: DNS, Firewall 1, Firewall 2, Router 1, Router 2, VPN and Orchestrator Server.

Identify the problematic device(s). Part 2:

Identify the correct options to provide adequate configuration for hybrid cloud architecture.

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

Part 1:

Cloud Hybrid Network Diagram

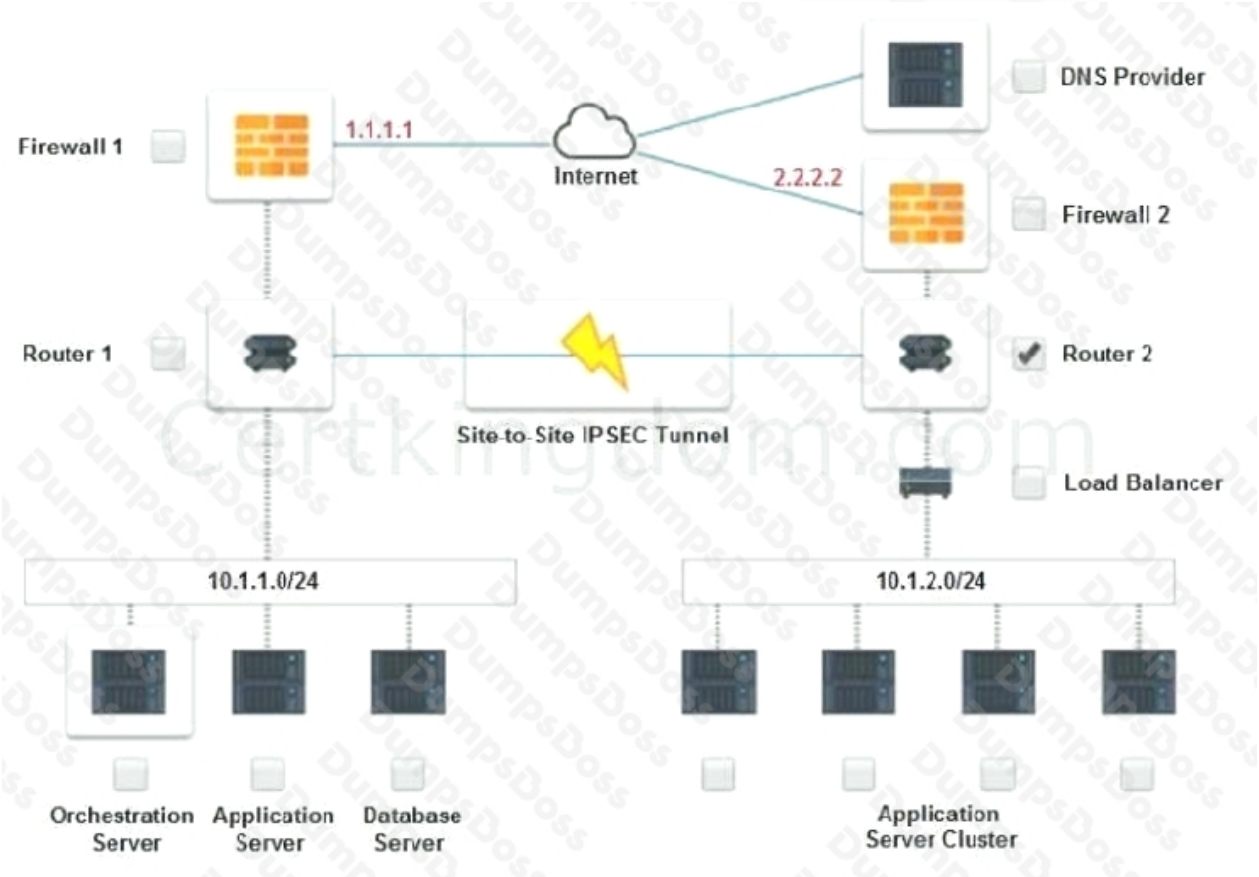

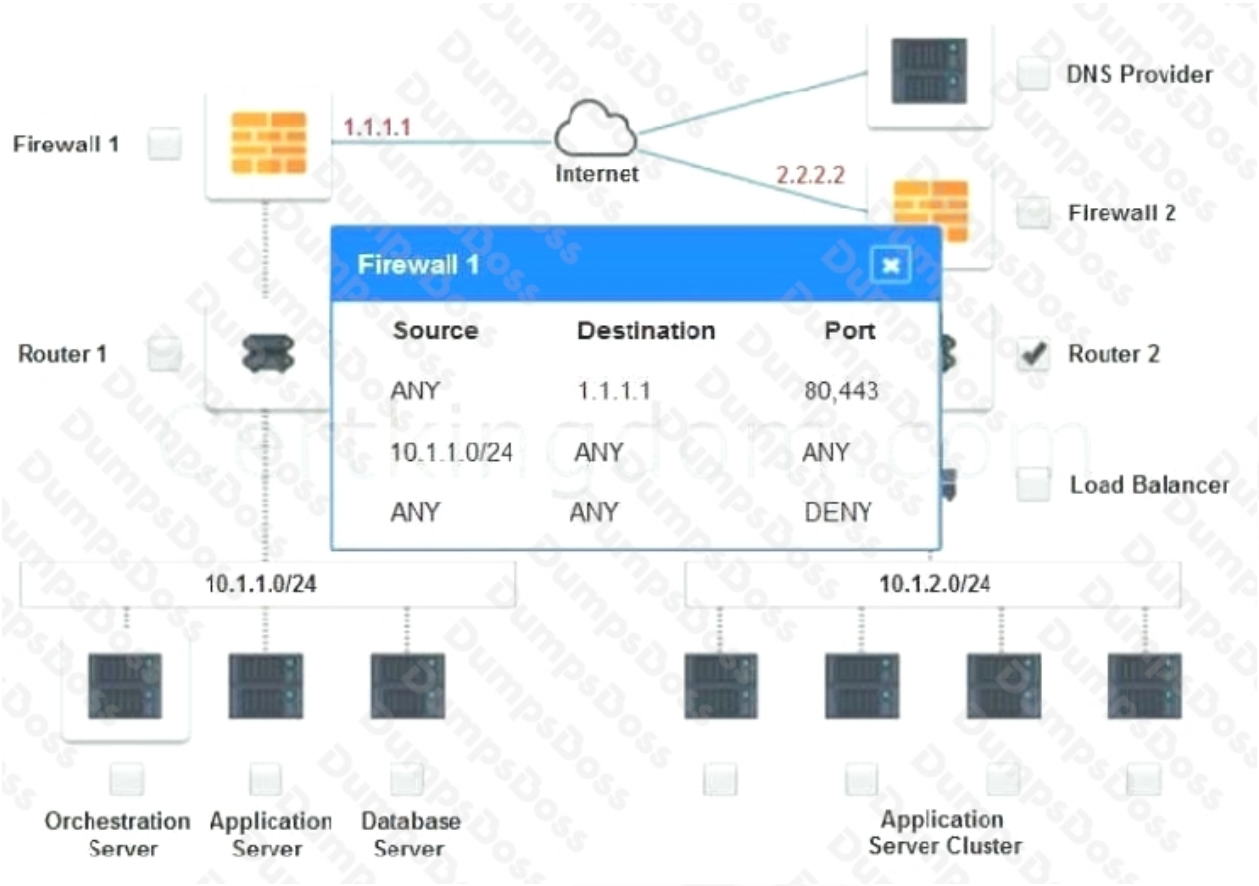
DNS Provider
Firewall 1
1.1.1.1
Internet
2.2.2.2
Firewall 2
Firewall 1
Source
Destination
Port
ANY
1.1.1.1
80,443
10.1.1.0/24
ANY
ANY
ANY
ANY
DENY
Router 1
Router 2
Load Balancer
10.1.1.0/24
10.1.2.0/24
Orchestration Server
Application Server
Database Server
Application Server Cluster

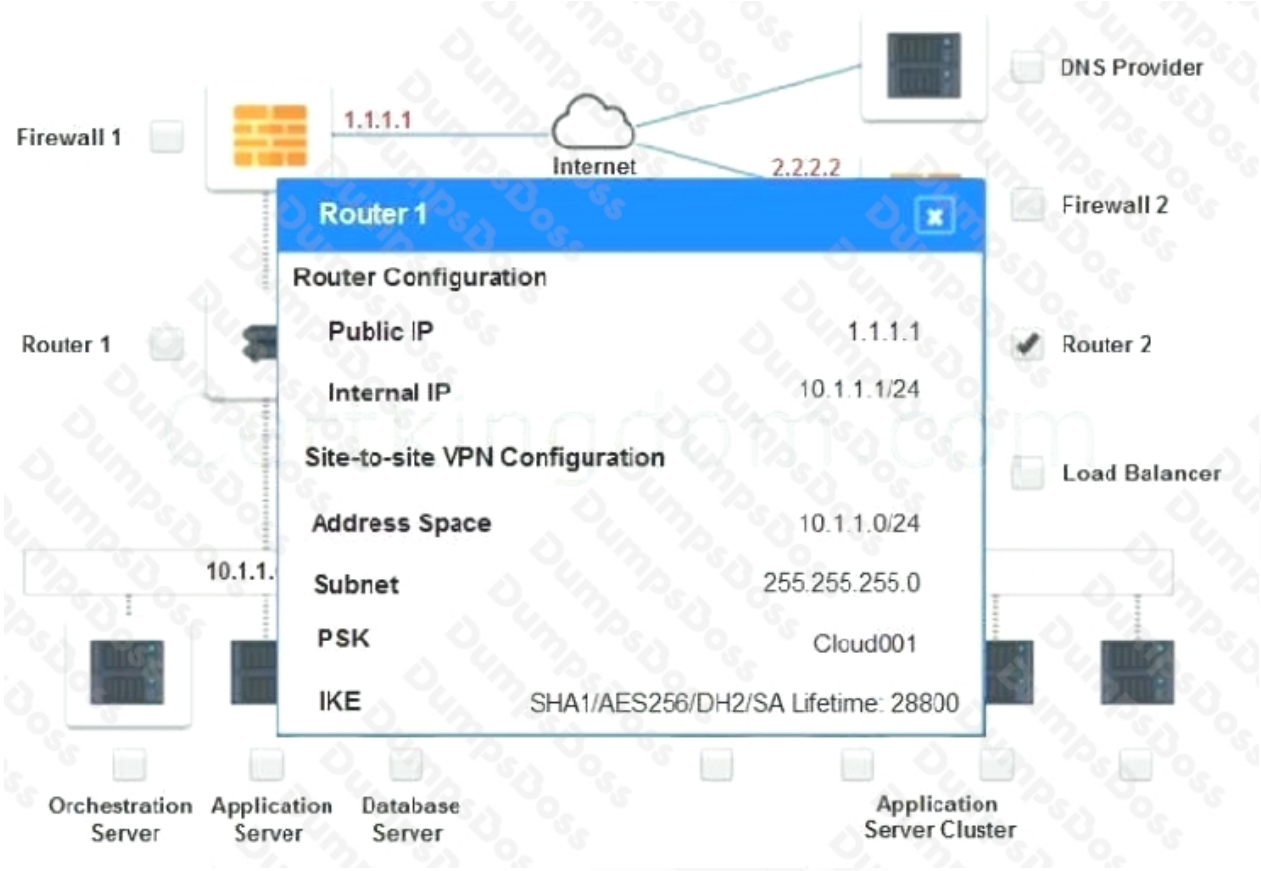
DNS Provider
Firewall 1
1.1.1.1
Internet
2.2.2.2
Router 1
Router Configuration
Public IP
1.1.1.1
Internal IP
10.1.1.1/24
Site-to-site VPN Configuration
Address Space
10.1.1.0/24
Subnet
255.255.255.0
PSK
Cloud001
IKE
SHA1/AES256/DH2/SA Lifetime: 28800
Firewall 2
Router 1
Router 2
Load Balancer
10.1.1.
Orchestration Server
Application Server
Database Server
Application Server Cluster

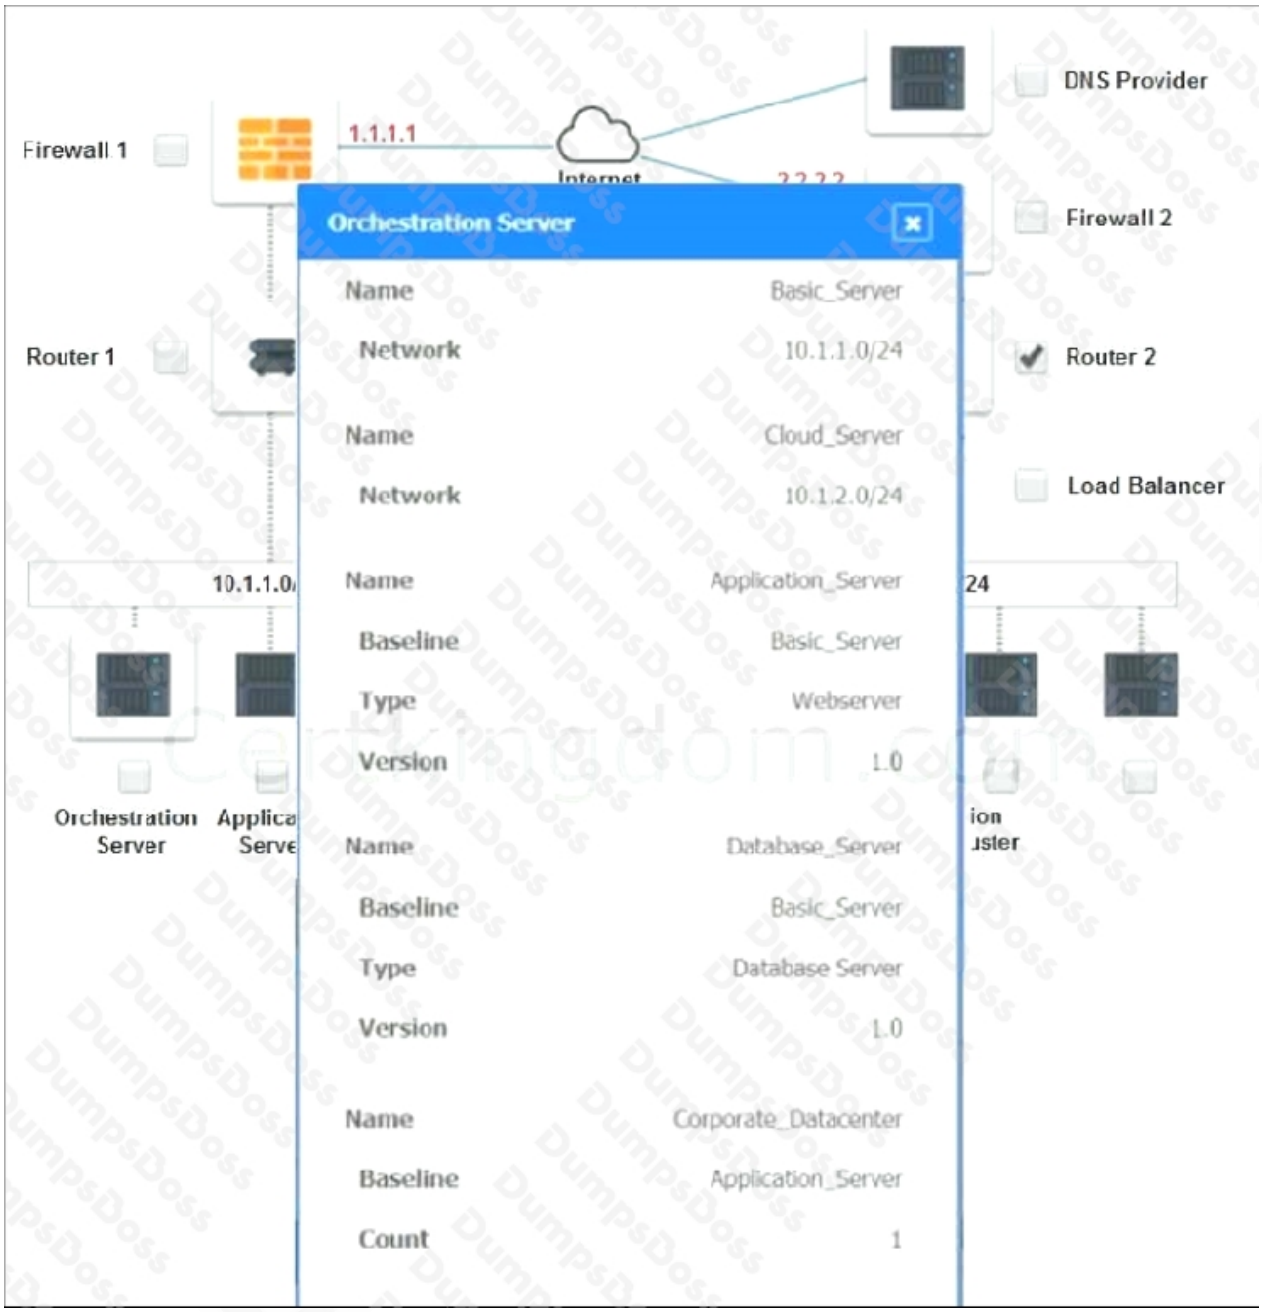
DNS Provider
Firewall 1
1.1.1.1
Internet
2.2.2.2
Firewall 2
Orchestration Server
Name
Basic_Server
Network
10.1.1.0/24
Router 1
Router 2
Name
Cloud_Server
Network
10.1.2.0/24
Load Balancer
10.1.1.0/
Name
Application_Server
24
Baseline
Basic_Server
Type
Webserver
Version
1.0
Orchestration Server
Applica
Serve
ion
uster
Name
Database_Server
Baseline
Basic_Server
Type
Database Server
Version
1.0
Name
Corporate_Datacenter
Baseline
Application_Server
Count
1

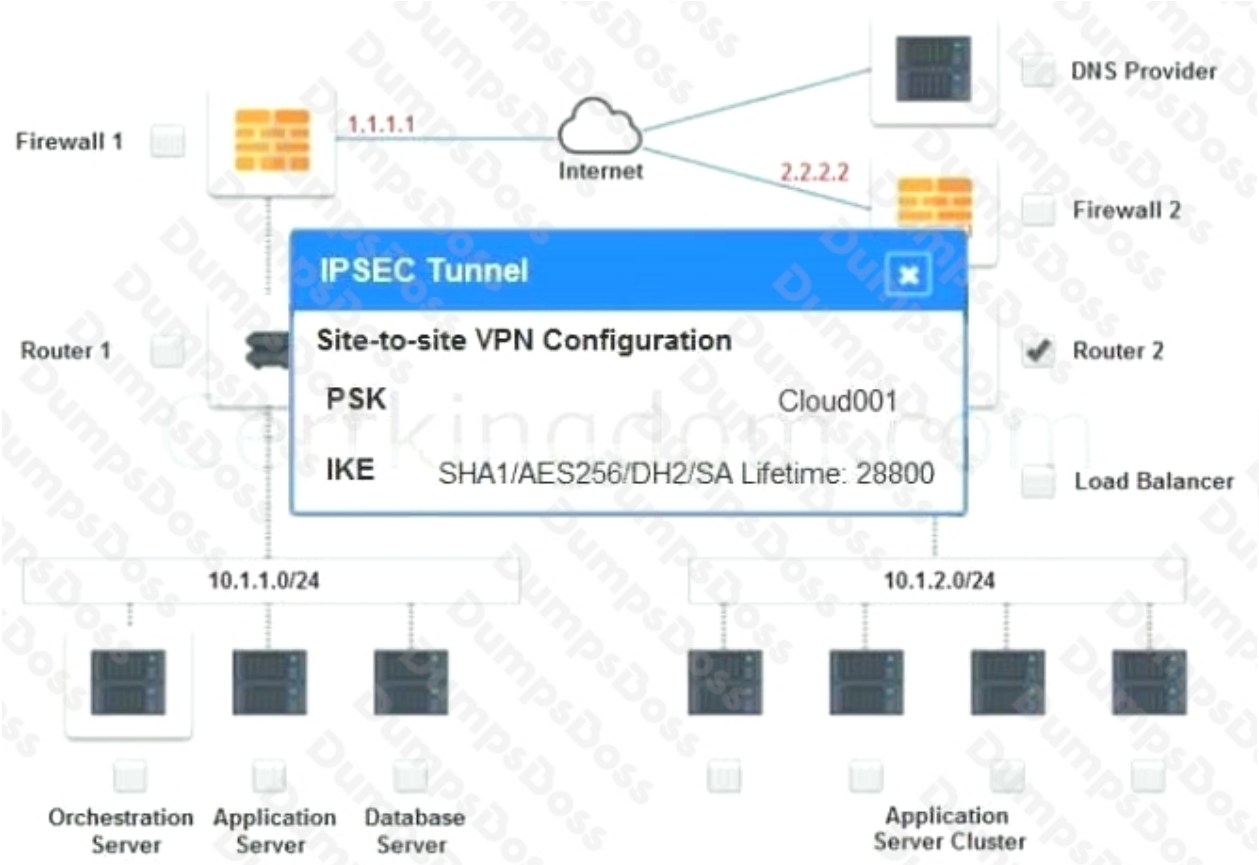
DNS Provider
Firewall 1
1.1.1.1
Internet
2.2.2.2
Firewall 2
IPSEC Tunnel
Site-to-site VPN Configuration
PSK
Cloud001
IKE
SHA1/AES256/DH2/SA Lifetime: 28800
Router 1
Router 2
Load Balancer
10.1.1.0/24
10.1.2.0/24
Orchestration Server
Application Server
Database Server
Application Server Cluster

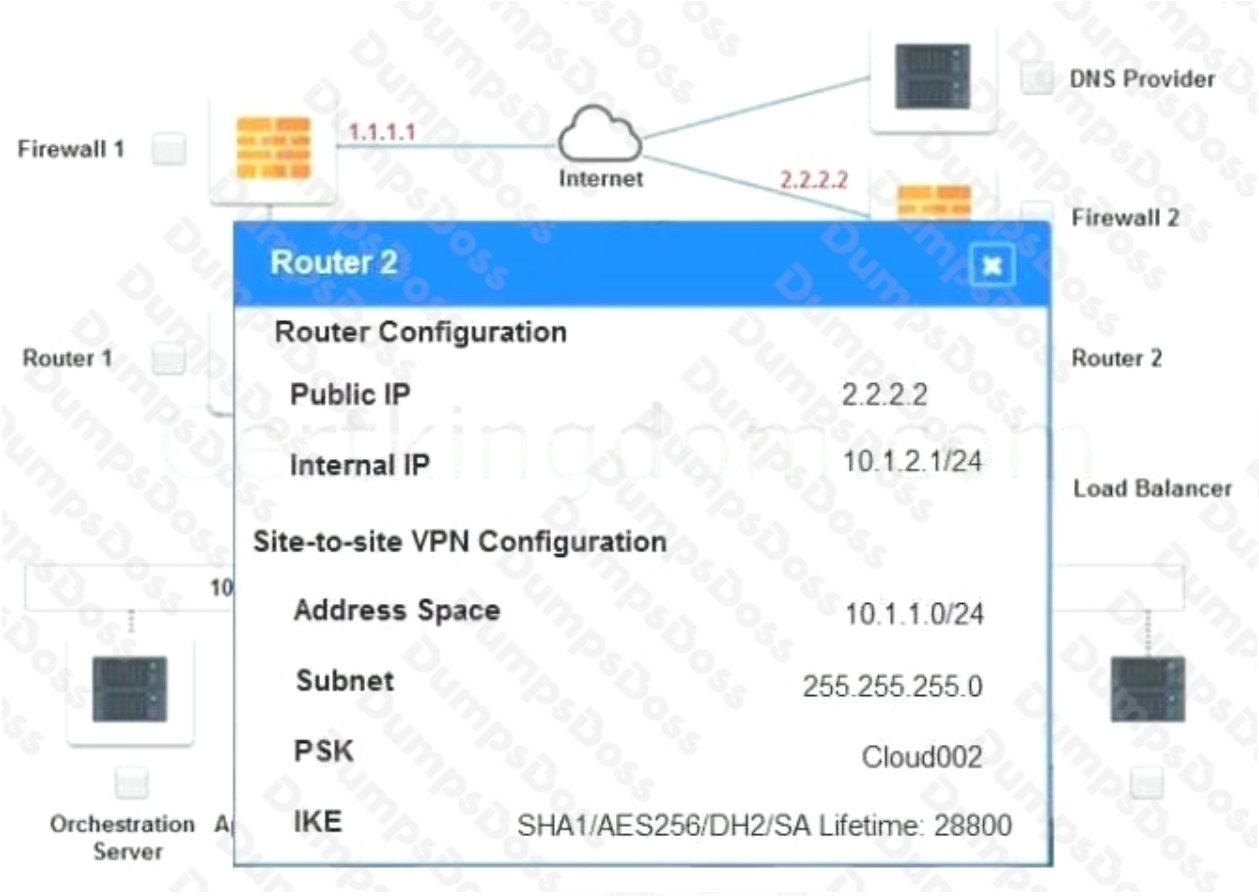
DNS Provider
Firewall 1
1.1.1.1
Internet
2.2.2.2
Firewall 2
Router 2
Router Configuration
Public IP
2.2.2.2
Internal IP
10.1.2.1/24
Router 1
Router 2
Load Balancer
Site-to-site VPN Configuration
Address Space
10.1.1.0/24
Subnet
255.255.255.0
PSK
Cloud002
IKE
SHA1/AES256/DH2/SA Lifetime: 28800
Orchestration Server

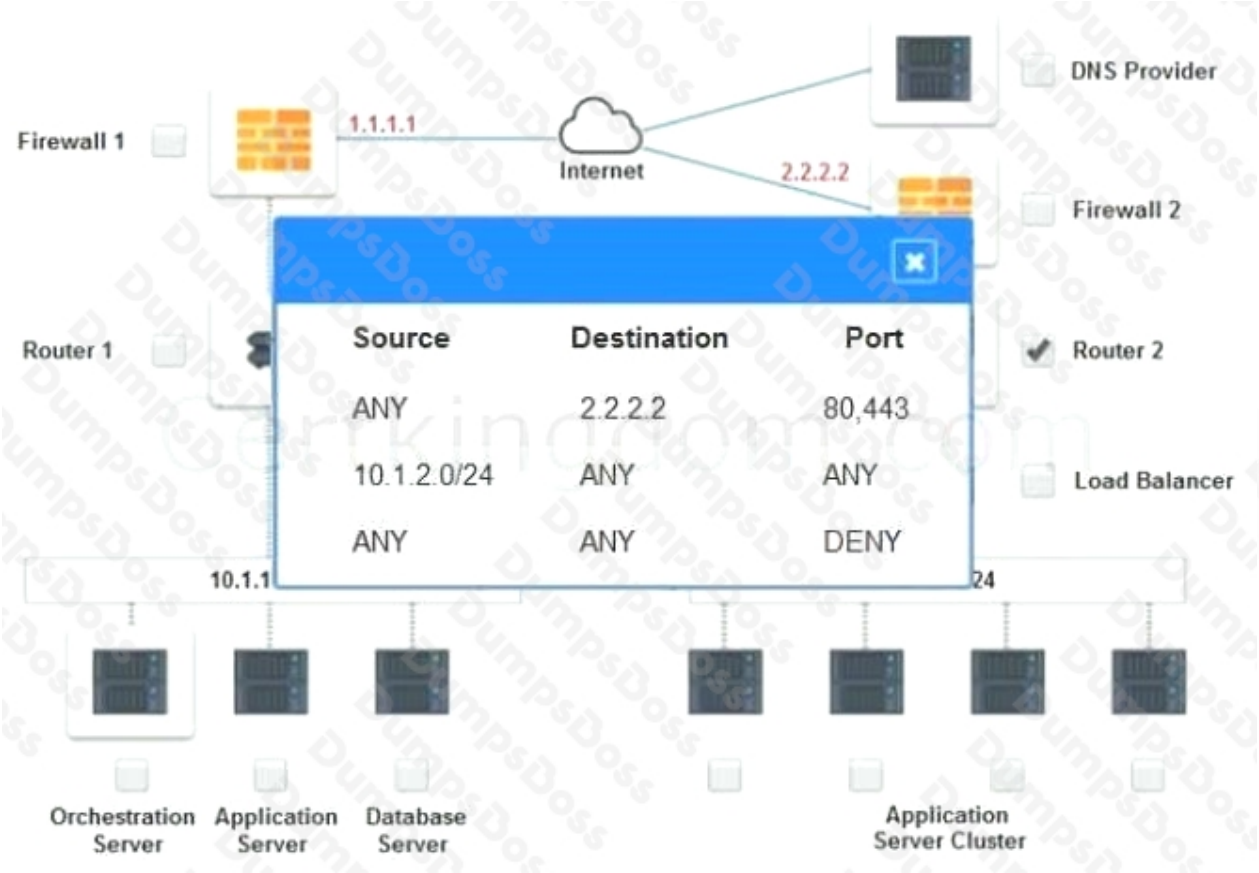
DNS Provider
Firewall 1
1.1.1.1
Internet
2.2.2.2
Firewall 2
Router 1
Router 2
Load Balancer
Source
Destination
Port
ANY
2.2.2.2
80,443
10.1.2.0/24
ANY
ANY
ANY
ANY
DENY
10.1.1
24
Orchestration Server
Application Server
Database Server
Application Server Cluster

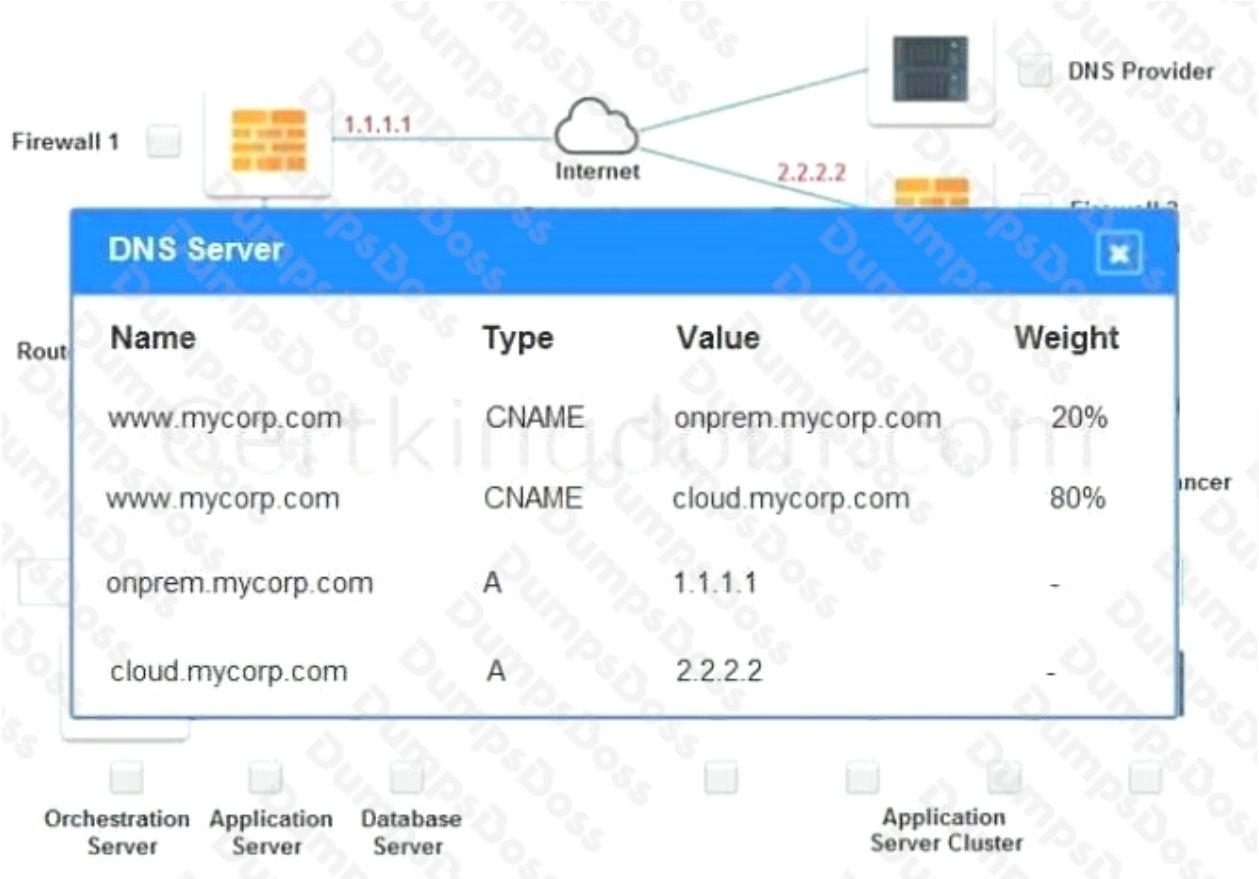

Part 2:

Only select a maximum of TWO options from the multiple choice question

- ☐ Deploy a Replica of the Database Server in the Cloud Provider.
- ☐ Update the PSK (Pre-shared key) in Router 2.
- ☐ Update the A record on the DNS from 2.2.2.2 to 1.1.1.1.
- ☐ Promote deny AI to allow All in Firewall 1 and Firewall 2.
- ☐ Change the Address Space on Router 2.
- ☐ Change internal IP Address of Router 1.
- ☐ Reverse the Weight property in the two CNAME records on the DNS.
- ☐ Add the Application Server at on-premises to the Load Balancer.

ANSWER: Seeexplanationbelow.

Explanation:

Part 1: Router 2

The problematic device is Router 2, which has an incorrect configuration for the IPSec tunnel. The IPSec tunnel is a secure connection between the on-premises datacenter and the cloud provider, which allows the traffic to flow between the two networks. The IPSec tunnel requires both endpoints to have matching parameters, such as the IP addresses, the pre-shared key (PSK), the encryption and authentication algorithms, and the security associations (SAs) .

According to the network diagram and the configuration files, Router 2 has a different PSK and a different address space than Router 1. Router 2 has a PSK of œ1234567890 , while Router 1 has a PSK of œ0987654321 . Router 2 has an address space of 10.0.0.0, while Router 1 has an address space of 192.168.0.0. These mismatches prevent the IPSec tunnel from establishing and encrypting the traffic between the two networks.

The other devices do not have any obvious errors in their configuration. The DNS provider has two CNAME records that point to the application servers in the cloud provider, with different weights to balance the load. The firewall rules allow the traffic from and to the application servers on port 80 and port 443, as well as the traffic from and to the VPN server on port 500 and port 4500. The orchestration server has a script that installs and configures the application servers in the cloud provider, using the DHCP server to assign IP addresses.

Part 2:

The correct options to provide adequate configuration for hybrid cloud architecture are:

Update the PSK in Router 2.

Change the address space on Router 2.

These options will fix the IPSec tunnel configuration and allow the traffic to flow between the onpremises datacenter and the cloud provider. The PSK should match the one on Router 1, which is œ0987654321 . The address space should also match the one on Router 1, which is 192.168.0.0.

B. Update the PSK (Pre-shared key in Router2) E. Change the Address Space on Router2

QUESTION NO: 11

A cloud engineer was deploying the company's payment processing application, but it failed with the

following error log:

ERFOR:root: Transaction failed http 429 response, please try again Which of the following are the most likely causes for this error? (Select two).

A. API throttling

B. API gateway outage

C. Web server outage

D. Oversubscription

E. Unauthorized access

F. Insufficient quota

ANSWER: A F

Explanation:

An HTTP 429 error means “Too Many Requests.” In plain terms, the app is hitting an API endpoint more often than it’s allowed to, so the service starts rejecting calls to protect itself. That’s why **API throttling** is a top suspect—rate limits are commonly enforced by cloud APIs and API gateways, and once you cross that limit you’ll see 429 responses.

The other common reason is **insufficient quota**. Even if each burst of requests is within a per-second limit, you can still run out of allowed requests for the hour/day/month (depending on the provider). When the quota is exhausted, providers may also respond with 429 until the quota resets or you request a higher limit.

Issues like outages usually show up as 5xx errors, and unauthorized access is typically 401/403, not 429. Good next steps are checking the provider’s rate-limit headers, reviewing usage/quota dashboards, and adding backoff/retry logic.

References: https://developer.mozilla.org/en-US/docs/Web/HTTP/Status/429 and https://cloud.google.com/apis/design/errors

QUESTION NO: 12

A cloud networking engineer is troubleshooting the corporate office's network configuration. Employees in the IT and operations departments are unable to resolve IP addresses on all devices, and the IT department cannot establish a connection to other departments' subnets. The engineer identifies the following configuration currently in place to support the office network:

Subnet	Department	Employees
10.1.20.1/24	Finance	50
10.1.30.1/24	IT	90
10.1.40.1/24	Legal	30
10.1.50.1/24	Operations	100

Each employee needs to connect to the network with a maximum of three hosts. Each subnet must be segregated, but the IT department must have the ability to communicate with all subnets. Which of the following meet the IP addressing and routing requirements? (Select two).

A. Modifying the subnet mask to 255 255 254.0 for IT and operations departments

B. Configuring static routing to allow access from each subnet to 10.1.40.1

C. Modifying the BYOD policy to reduce the volume of devices that are allowed to connect to the corporate network

D. Configuring static routing to allow access from 10.1.30.1 to each subnet

E. Combining the subnets and increasing the allocation of IP addresses available to support three hosts for each employee

F. Modifying the subnet mask to 255.255.255.128 for the IT and operations departments

G. Modifying the subnet mask to 255.255.254.0 (/23) for IT and Operations to support up to ~510 usable IPs per subnet

ANSWER: A D G

Explanation:

IT needs to talk to every other subnet, but the departments still have to stay separated. The clean way to do that is routing: you add routes so the IT subnet (10.1.30.0/24) knows how to reach Finance, Legal, and Operations (and the return routes exist on the other side, usually via a router/L3 switch). That's exactly what static routing from 10.1.30.1 to each subnet is meant to solve.

On the addressing side, Operations has 100 employees and IT has 90. If each employee may need up to three hosts, those networks need room for up to ~300 and ~270 IPs. A /24 only gives 254 usable addresses, so it's too small. Changing those two to a /25 (255.255.255.128) would actually make the problem worse (only 126 usable). So while the routing choice is right, the subnet mask choice in the options doesn't truly meet the "three hosts per employee" requirement. A more correct mask would be /23 (255.255.254.0) for IT and Operations.

References: https://www.iana.org/help/netmask and https://en.wikipedia.org/wiki/Classless_Inter-Domain_Routing

QUESTION NO: 13

A security analyst confirms a zero-day vulnerability was exploited by hackers who gained access to

confidential customer data and installed ransomware on the server Which of the following steps should the security analyst take? (Select two).

A. Contact the customers to inform them about the data breach.

B. Contact the hackers to negotiate payment lo unlock the server.

C. Send a global communication to inform all impacted users.

D. Inform the management and legal teams about the data breach

E. Delete confidential data used on other servers that might be compromised.

F. Modify the firewall rules to block the IP addresses and update the ports.

ANSWER: A D

Explanation:

When you've got a confirmed breach involving customer data plus ransomware, two of the first "must-do" actions are notification and escalation. Informing management and legal right away is key because they'll drive the incident response process, help preserve evidence, and make sure you follow breach notification laws and contracts. Legal also helps control what gets communicated and when, so the company doesn't accidentally create more risk.

You should also notify affected customers (typically through an approved breach-notification process). Since confidential data was accessed, customers need to know what happened and what steps they should take next (like monitoring accounts or changing passwords). This is also commonly required by regulations depending on where the customers live.

Options like negotiating with attackers or making a "global" announcement aren't the right first moves for a security analyst. Blocking IPs or changing ports might help later, but it doesn't replace proper incident handling and required notifications.

References: https://www.cisa.gov/stopransomware/ransomware-guide and https://www.ftc.gov/business-guidance/resources/data-breach-response-guide-business

QUESTION NO: 14

A CI/CD pipeline is used to deploy VMs to an IaaS environment. Which of the following can be used to harden the operating system once the VM is running?

A. Docker

B. Kubernetes

C. Git

D. Ansible

ANSWER: D

Explanation:

Ansible is the best fit here because it's a configuration management and automation tool. After a VM boots, you can use Ansible playbooks to apply OS hardening steps consistently—things like disabling unused services, enforcing secure SSH settings, setting file permissions, applying patches, and pushing security baselines.

The other options don't really do OS hardening for running VMs. Docker and Kubernetes are mainly for building and orchestrating containers, not locking down the underlying VM operating system. Git is a version control system; it can store your hardening scripts, but it doesn't actually apply changes to the VM by itself.

References: https://docs.ansible.com/ansible/latest/getting_started/index.html and https://www.redhat.com/en/technologies/management/ansible/what-is-ansible

QUESTION NO: 15

Following a ransomware attack, the legal department at a company instructs the IT administrator to store the data from the affected virtual machines for a minimum of one year. Which of the following is this an example of?

A. Recoverability

B. Retention

C. Encryption

D. Integrity

ANSWER: B

Explanation:

This is a classic example of a data retention requirement. The key clue is "store the data… for a minimum of one year." That's not about fixing systems (recoverability), protecting data with cryptography (encryption), or making sure data wasn't altered (integrity). It's simply a rule that says how long the company must keep certain data before it can be deleted.

After an incident like ransomware, legal teams often place a "legal hold" or set a retention period so evidence is preserved for investigations, insurance claims, audits, or possible lawsuits. In practice, IT might snapshot the affected VMs, export logs, and archive disk images in immutable storage—specifically to meet that time-based requirement.

References: https://www.nist.gov/privacy-framework/nist-privacy-framework/glossary/data-retention and https://www.cisa.gov/news-events/news/implementing-legal-holds-and-retention (general guidance on preserving incident-related data).