1Z0-997-22 Q&As

Oracle Cloud Infrastructure 2022 Architect Professional

Pass Oracle 1Z0-997-22 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

https://www.leads4pass.com/1z0-997-22.html

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Oracle Official Exam Center

QUESTION 1

As a part of migration exercise for an existing on premises application to Oracle Cloud Infrastructure (OCT), yon ore required to transfer a 7 TB file to OCI Object Storage. You have decided to upload functionality of Object Storage. Which two statements are true?

A. Active multipart upload can be checked by listing all parts that have been uploaded, however It Is not possible to list information for individual object part in an active multipart upload

B. It is possible to spill this fileInto multiple parts using the APIs provided by Object Storage.

C. It is possible to split this file into multiple parts using rclone tool provided by Object Storage.

D. After initiating a multipart upload by making a CreateMultlPartUpload RESI API Call, the upload remains active until you explicitly commit it or abort.

E. Contiguous numbers need to be assigned for each part so that Object Storage constructs the object by ordering, part numbers in ascending order

Correct Answer: AD

You can check on an active multipart upload by listing all parts that have been uploaded. (You cannot list information for an individual object part in an active multipart upload.)

After you finish creating object parts, initiate a multipart upload by making a CreateMultipartUpload REST API call. Provide the object name and any object metadata. Object Storage responds with a unique upload ID that you must include in any requests related to this multipart upload. Object Storage also marks the upload as active. The upload remains active until you explicitly commit it or abort it.

QUESTION 2

You have designed and deployed your Autonomous Data Warehouse (ADW) such that it is accessible from your on-premises data center and servers running on both private and public networks in Oracle Cloud Infrastructure (OCI).

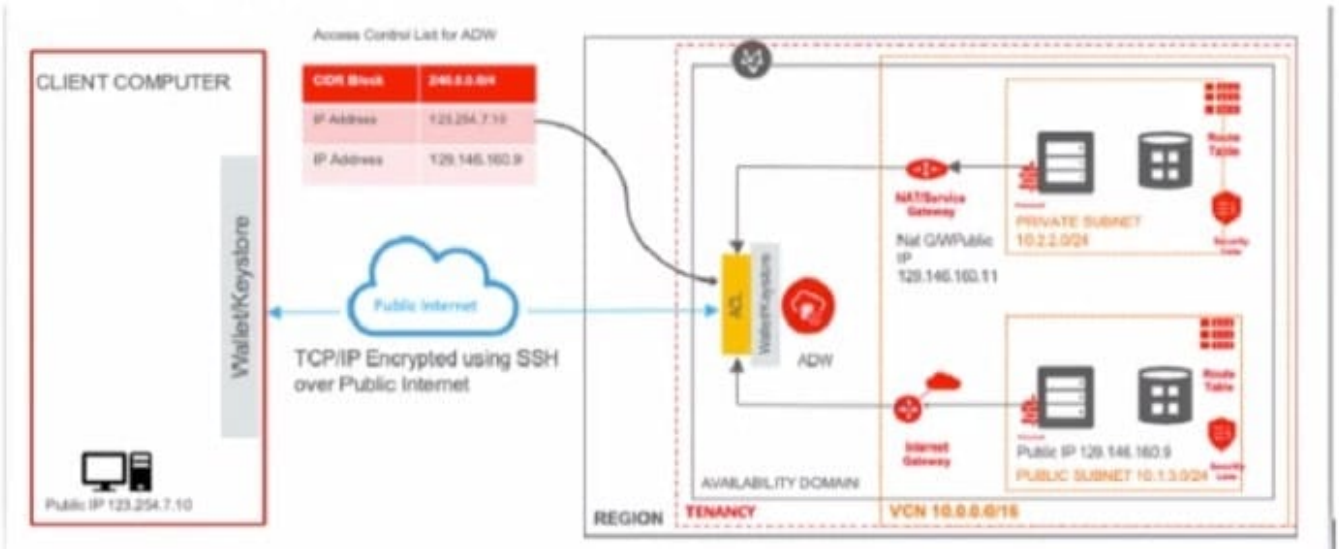

As you are testing the connectivity to your ADW database from the different access paths, you notice that the server

running on the private network is unable to connect to ADW.

Which two steps do you need to take to enable connectivity from the server on the private network to ADW? (Choose two.)

A. Add an entry in the Security List of the ADW allowing ingress traffic for C10R block 10.2.2.0/24

B. Add an entry in the route table (associated with the private subnet) with destination of 0.0.0.0/: target type of NAT Gateway, add a stateful egress rule to the security list (associated with the private subnet) with destination of 0.0.0.0./0 and for all IP protocols.

C. Add an entry in the access table list of ASW for CIDR block 10.2.2.0/24.

D. Add an entry in the route table (associated with the private subnet) with destination of 0.0.0.0./0; target type of internet Gateway, add a stateful egress in the security list (associated with the private subnet) with destination of 0.0.0.0/0 and for all IP protocols.

E. Add an entry in the access control list of ADW for IP address 129.146.160.11

Correct Answer: BE

There are 3 connections to ADW 1- Connecting to (ADW) from Public Internet 2- Connecting to ADW (via NAT or Service Gateway) from a server running on a private subnet in OCI (in the same tenancy) 3- Connecting to ADW (via internet Gateway) from a server running on a public subnet in OCI (in the same tenancy

QUESTION 3

A cloud consultant is working on implementation project on OCI. As part of the compliance requirements, the objects placed in object storage should be automatically archived first and then deleted. He is testing a Lifecycle Policy on Object

Storage and created a policy as below:

[{ "name": "Archive_doc", "action": "ARCHIVE", "objectNameFilter": { "inclusionPrefixes":

"doc"] },

"timeAmount": 5, "timeunit": "DAYS", "isEnabled": true },

{ "name": "Delete_doc", "action": "DELETE", "objectNameFilter": "inclusionPrefixes": ["doc"]

1."timeAmount": 5, "timeunit": "DAYS", "isEnabled": true }

What will happen after this policy is applied?

A. All objects with names starting with "doc" will be deleted after 5 days of object creation

B. All the objects having file extension ".doc" will be archived for 5 days and will be deleted 10 days after object creation

C. All the objects having file extension ".doc" will be archived 5 days after object creation

D. All the objects with names starting with "doc" will be archived 5 days after object creation and will be deleted 5 days after archival

Correct Answer: A

Object Lifecycle Management works by defining rules that instruct Object Storage to archive or delete objects on your behalf within a given bucket. A bucket\\'s lifecycle rules are collectively known as an object lifecycle policy. You can use a rule to either archive or delete objects and specify the number of days until the specified action is taken.

A rule that deletes an object always takes priority over a rule that would archive that same object.

QUESTION 4

You are advising the database administrator responsible for managing non-production environment for Oracle Autonomous Database running on Oracle Cloud Infrastructure. You need to help the database administrator ensure that the non-production environments have a copy of the current data from the production environment in a manner that is most time-efficient.

Which method should you recommend? (Choose the best answer.)

A. Take a full database backup of the production Autonomous database and create the non-production database from it.

B. Create a metadata clone of the production Autonomous Database and create the non- production database from it.

C. Create a full clone of the production Autonomous Database and create the non- production database from it.

D. Take a Data Pump export of the production Autonomous database and import into the non-production database.

Correct Answer: C

Explanation: https://www.oracle.com/database/technologies/datawarehouse-bigdata/adb- faqs.html

QUESTION 5

You notice that a majority of your Oracle Cloud Infrastructure (OCI) resources like compute instances, block volumes, and load balancers are not tagged. You have received a mandate from your CIO to add a predefined set of tags to identify owners for respective OCI resources. E.g. if Chris and Larry each create compute instances in a compartment, the instances that Chris creates include tags that contain his name as the value, while the instances that Larry creates have his name.

Which option is the simplest way to implement this new tagging requirement?

A. Create a default tag for each compartment, which ensure that appropriate tags are applied at the time of resource creation.

B. Create an OCI Identity and Access Management policy requiring users to tag resources with their user name.

C. Create an OCI Identity and Access Management policy to automatically tag a resource with the user name.

D. Create tag variables to automatically tag a resource with the user name.

Correct Answer: D

QUESTION 6

You have deployed a multi-tier application with multiple compute instances in Oracle Cloud Infrastructure. You want to back up these volumes and have decided to use Volume Group\\'s feature. The Block volume and Compute instances exist in different compartments within your tenancy.

Periodically. a few child compartments are moved under different parent compartments, and you notice that sometimes volume group backup fails.

What could be the cause?

A. You are exceeding your volume group backup quota configured.

B. You have the same block volume attached to multiple compute instances; if these compute instances are in different compartments then all concerned compartments must be moved at the same time.

C. Compute instance with multiple block volumes attached cannot move when a compartment is moved.

D. The Identity and Access Management policy allowing backup failed to move when the compartment was moved.

Correct Answer: D

You can move a compartment to a different parent compartment within the same tenancy. When you move a compartment, all its contents (subcompartments and resources) are moved with it. Moving a compartment has implications for the

contents. After you move a compartment to a new parent compartment, the access policies of the new parent take effect and the policies of the previous parent no longer apply. Before you move a compartment, ensure that:

You are aware of the policies that govern access to the compartment in its current position. You are aware of the polices in the new parent compartment that will take effect when you move the compartment.

In some cases, when moving nested compartments with policies that specify the hierarchy, the polices are automatically updated to ensure consistency.

QUESTION 7

Your Oracle database is deployed on-premises and has produced 100 TB database backup locally. You have a disaster recovery plan that requires you to create redundant database backups in Oracle Cloud Infrastructure (OCI).

Once the initial backup is completed, the backup must be available for retrieval in less than 30 minutes to support the Recovery Time Objective (RTO) of your solution.

Which is the most cost effective option to meet these requirements?

A. Setup an IPsec VPNConnect between on-premises data center and OCI. Then to use OCI CLI command to upload database backups to OCI Object Storage Archive tier as the final destination.

B. Use OCI Storage Gateway to transfer the backup files to OCI Object Storage Archive tier as the final destination.

C. Setup a FastConnect connection between on-premises data center and OCI. Then to use OCI CLI command to upload database backups to OCI Object Storage Standard tier as the final destination.

D. Use OCI Storage Gateway to transfer the backup files to OCI Object Storage Standard tier as the final destination.

Correct Answer: D

QUESTION 8

All three Data Guard Configuration are fully supported on Oracle Cloud infrastructure (OCI). You want to deploy a maximum availability architecture (MAA) for database workload.

Which option should you consider while designing your Data Guard configuration to ensure best RTO and PRO without causing any data loss?

A. Configure "Maximum Protection" mode which provides zero data loss If the primary database fails.

B. Configure "Maximum Performance" mode In SYNC mode between two availability domains (same region) which provides, the highest level of data protection that is possible without affecting the performance of the primary database.

C. Configure \\'\\'Maximum Scalability" mode which provides the highest level of scalability without compromising the availability of the primary database.

D. Configure \\'\\'Maximum Availability" mode in SYNC mode between two availability domains (same region), and use the Maximum Availability mode in SYNC mode between two regions.

Correct Answer: D

Explanation: https://docs.cloud.oracle.com/enus/iaas/Content/Resources/Assets/whitepapers/best-practices-for-dr-on-oci.pdf

All three Data Guard configurations are fully supported on Oracle Cloud Infrastructure. However, because of a high risk of production outage, we don\\'t recommend using the maximum protection mode for your Data Guard configuration. We recommend using the maximum availability mode in SYNC mode between two availability domains (same region), and using the maximum availability mode in ASYNC mode between two regions. This architecture provides you the best RTO and RPO without causing any data loss. We recommend building this architecture in daisy-chain mode: the primary database ships redo logs to the first standby database in another availability domain in SYNC mode, and then the first standby database ships the redo logs to another region in ASYNC mode. This method ensures that your primary database is not doing the double work of shipping redo logs, which can cause performance impact on a production workload.

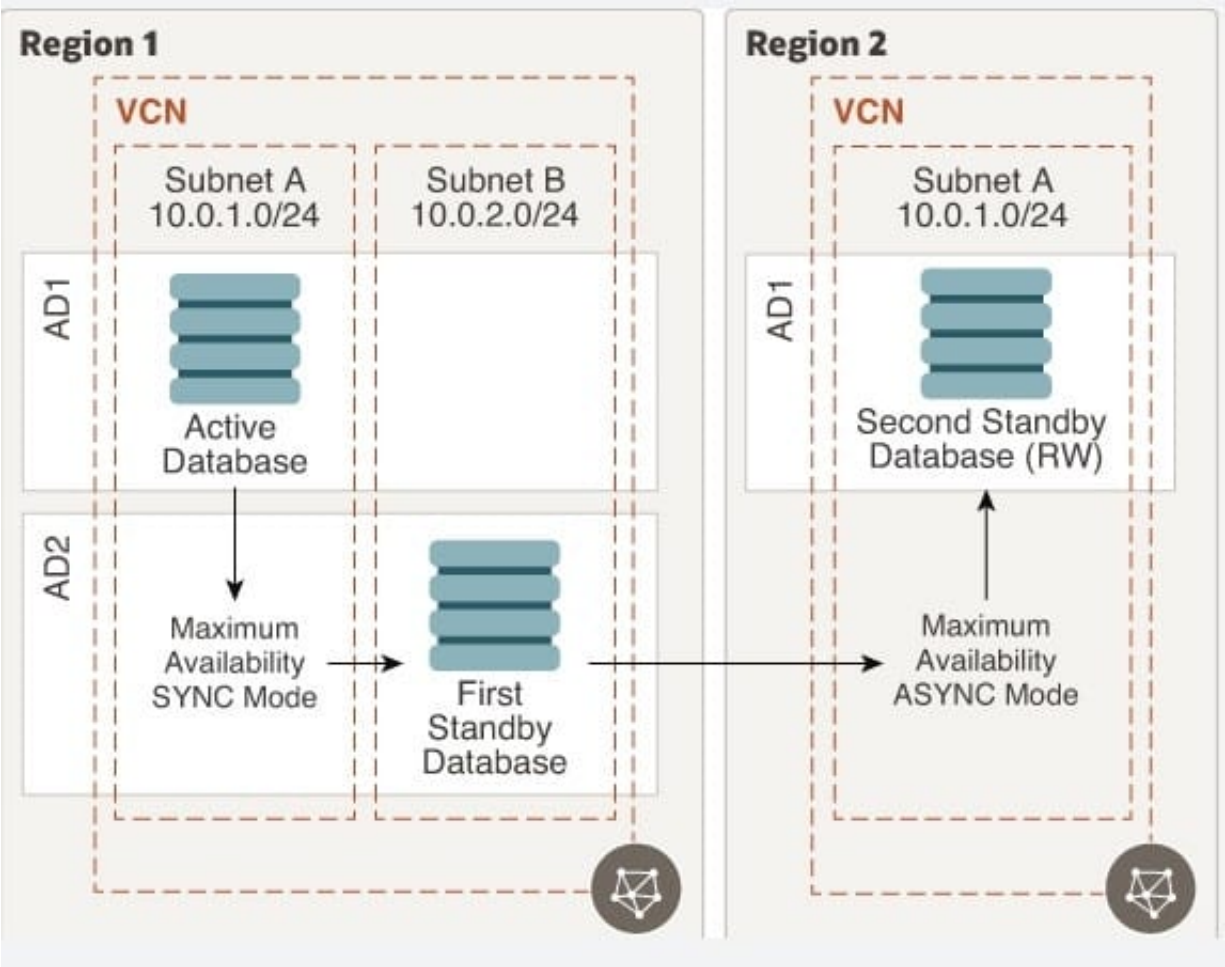

This configuration offers the following benefits: No data loss within a region. No overhead on the production database to maintain standbys in another region. Option to configure lagging on the DR site if needed for business reasons. Option to configure multiple standbys in different regions without any additional overhead on the production database. A typical use case is a CDN application Bottom of Form

QUESTION 9

An E-commerce company which sells computers, tablets, and other electronics items has recently decided to move all of their on-premises infrastructure to Oracle Cloud Infrastructure (OCI). One of their on-premises application is running on an NGINX server and the Oracle Database is running in a 2 node Oracle Real Application Clusters (RAC) configuration.

They cannot afford to have any application down time when they do the migration.

What is an effective mechanism to migrate the customer application to OCI and set up regular automated backups?

A. Launch a compute instance and run an NGINX server to host the application. Deploy a 2 node VM DB Systems with Oracle RAC enabled. Import the on-premises database to OCI VM DB Systems using Oracle Data Pump and then enable automatic backups.

B. Launch a compute instance for both the NGINX application server and the database server. Attach block volumes on the database server compute instance and enable backup policy to backup the block volumes.

C. Launch a compute instance and run an NGINX server to host the application. Deploy Exadata Quarter Rack, enable automatic backups and import the database using Oracle Data Pump.

D. Launch a compute instance and run an NGINX server to host the application. Deploy a 2 node VM DB Systems with Oracle RAC enabled. Setup Oracle GoldenGate to synchronize data from their on-premises database to OCIVM Database. Export and Import the on- premises database to OCIVM DB Systems using Oracle Data Pump, apply the GoldenGate trail files to sync up the OCI database with the on-premises database. Enable automatic backups for the OCIVM database and then cutoverthe application from on-premises to OCI.

Correct Answer: D

QUESTION 10

Which three options are available to migrate an Oracle database 12.x from an on-premises environment to Oracle Cloud Infrastructure (OCI)?

A. Leverage OCI Storage Gateway asynchronous database migration option.

B. Use Oracle Data Pump Export/Import to migrate the database.

C. Configure RMAN cross-platform transportable tablespace backup sets.

D. Setup OCI schema and data transfer tool with Bare Metal DB Systems as the target.

E. Create a backup of your on-premises database In OCI DB Systems.

Correct Answer: BCE

Explanation: https://docs.cloud.oracle.com/en-us/iaas/Content/Database/Tasks/mig- onprembackup.htm

QUESTION 11

A data analytics company has been building its next-generation big data and analytics platform on Oracle Cloud Infrastructure (OCI) in the US East (Ashburn) region. They need a storage service that provides the scale and performance that their big data applications require such as high throughput to compute nodes coupled with low latency file operations.

In addition, they need to allow concurrent connections from multiple compute instances hosted in multiple Availability Domains and want to be able to quickly restore a previous version of the data in case of a need to roll back any major update.

Which option can they use to meet these requirements in the most cost-effective way?

A. Create a file system and mount target in the OCI File Storage service. Mount it into all the required compute instances. Take snapshots of the file system before each update.

B. Create block volume, attach it with read/write, shareable access type to all the required compute instances. Take a backup of the volume before each update.

C. Create an Object Storage bucket with object versioning enabled. Provision a compute instance to host the Storage Gateway and share the bucket via NFS, Mount the NFS into all the required compute instances.

D. Create a connection with the on-premises data center via FastConnect. Mount the shared NFS hosted on-premises.

Correct Answer: A

QUESTION 12

An insurance company is storing critical financial data in the OCI block volume. This volume is currently encrypted using oracle managed keys. Due to regulatory compliance, the customer wants to encrypt the data using the keys that they

can control and not the keys which are controlled by Oracle.

What of the following series of tasks are required to encrypt the block volume using customer managed keys?

A. Create a vault, import your master encryption key into the vault, generate data encryption key, assign data encryption key to the block volume

B. Create a master encryption key, create a data encryption key, decrypt the block volume using existing oracle managed keys, encrypt the block volume using the data encryption key

C. Create a vault, create a master encryption key in the vault, assign this master encryption key to the block volume

D. Create a master encryption key, create a new version of the encryption key, decrypt the block volume using existing oracle managed keys and encrypt using new version of the encryption key

Correct Answer: C

Oracle Cloud Infrastructure Vault lets you centrally manage the encryption keys that protect your data and the secret credentials that you use to securely access resources. You can use the Vault service to create and manage the following resources: Vaults Keys Secrets Vaults securely store master encryption keys and secrets that you might otherwise store in configuration files or in code. The Vault service lets you create vaults in your tenancy as containers for encryption keys and secrets. If needed, a virtual private vault provides you with a dedicated partition in a hardware security module (HSM), offering a level of storage isolation for encryption keys that\\'s effectively equivalent to a virtual independent HSM.

QUESTION 13

An online gaming application is deployed to multiple Availability Domains in the Oracle Cloud Infrastructure (OCI) us-ashburn-1 region. Considering the high volume of traffic that the gaming application handles, the company has hired you to ensure that the data stored by the application is scalable, highly available, and disaster resilient. In the event of failure, the Recovery Time Objective (RTO) and Recovery Point Objective (RPO) must be less than 2 hours.

Which Disaster Recovery strategy should be used to achieve the RTO and RPO requirements in the event of a system failure?

A. Configure hourly block volumes backups using the OCI Command Line Interface (CLI).

B. Create a user defined backup policy with a schedule of generating daily backups for block volumes.

C. Configure hourly block volumes backups through the OCI Storage Gateway service.

D. Create a user defined backup policy with a schedule of generating hourly backups for block volumes.

Correct Answer: A

QUESTION 14

You are building a highly available and fault tolerant web application deployment for your company. Similar application delayed by competitors experienced web site attack including DDoS which resulted in web server failing.

You have decided to use Oracle Web Application Firewall (WAF) to implement an architecture which will provide protection against such attacks and ensure additional configuration will you need to implement to make sure WAF is protecting my web application 24?.

Which additional configuration will you need to Implement to make sure WAF Is protecting my web application 24??

A. Configure auto scaling policy and it to WAF instance.

B. Configure Control Rules to send traffic to multiple web servers

C. Configure multiple origin servers

D. Configure new rules based on now vulnerabilities and mitigations

Correct Answer: C

Origin Management

An origin is an endpoint (typically an IP address) of the application protected by the WAF.

An origin can be

an Oracle Cloud Infrastructure load balancer public IP address. A load balancer IP address can be used for

high availability to an origin. Multiple origins can be defined, but only a single origin can be active for a WAF. You can set HTTP headers for outbound traffic from the WAF to the origin server. These name value pairs are then available to the

application. Oracle Cloud Infrastructure Web Application Firewall (WAF) is a cloud-based, Payment Card Industry (PCI) compliant, global security service that protects applications from malicious and unwanted internet traffic.

WAF can protect any internet facing endpoint, providing consistent rule enforcement across a customer\\'s applications. WAF provides you with the ability to create and manage rules for internet threats including Cross-Site Scripting (XSS),

SQL Injection and other OWASP- defined vulnerabilities. Unwanted bots can be mitigated while tactically allowed desirable bots to enter. Access rules can limit based on geography or the signature of the request.

Distributed Denial of Service (DDoS)

A DDoS attack is an often intentional attack that consumes an entity\\'s resources, usually using a large number of distributed sources. DDoS can be categorized into either Layer 7 or Layer 3/4 (L3/4)

A layer 7 DDoS attack is a DDoS attack that sends HTTP/S traffic to consume resources and hamper a website\\'s ability to delivery content or to harm the owner of the site. The Web Application Firewall (WAF)

service can protect layer 7 HTTP-based resources from layer 7 DDoS and other web application attack vectors.

QUESTION 15

Your customer has gone through a recent departmental re structure. As part of this change, they are organizing their

Oracle Cloud Infrastructure (OCI) compartment structure to align with the company\\'s new organizational structure.

They have made the following change:

Compartment x Is moved, and its parent compartment is now compartment c.

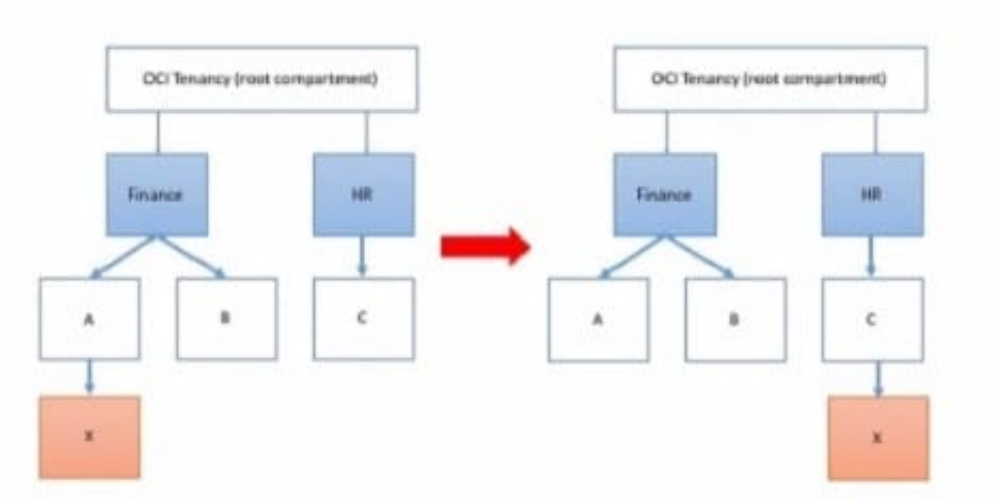

Policy defined in compartment A: Allow group networkadmins to manage subnets in compartment X Policy defined in root compartment: Allow group admins to read subnets in compartment Finance:A:X After you move the compartment, which two IAM policies would be required to ensure both groups retain the same permissions to compartment X that they had before? (Choose two.)

A. Define a policy in the root compartment as follows: Allow group admins to manage subnets in compartment Finance:A:X

B. Define a policy in compartment HR as follows: Allow group networkadmins to manage subnets in compartment C:X.

C. Define a policy in the root compartment as follows: Allow group admins to read subnets in compartment HR:C:X

D. Define a policy in compartment C as follows: Allow group networkadmins to read subnets in compartment X

Correct Answer: BC