Changing the Concept of 1z0-997-22 Exam Preparation 2022 [Q84-Q102

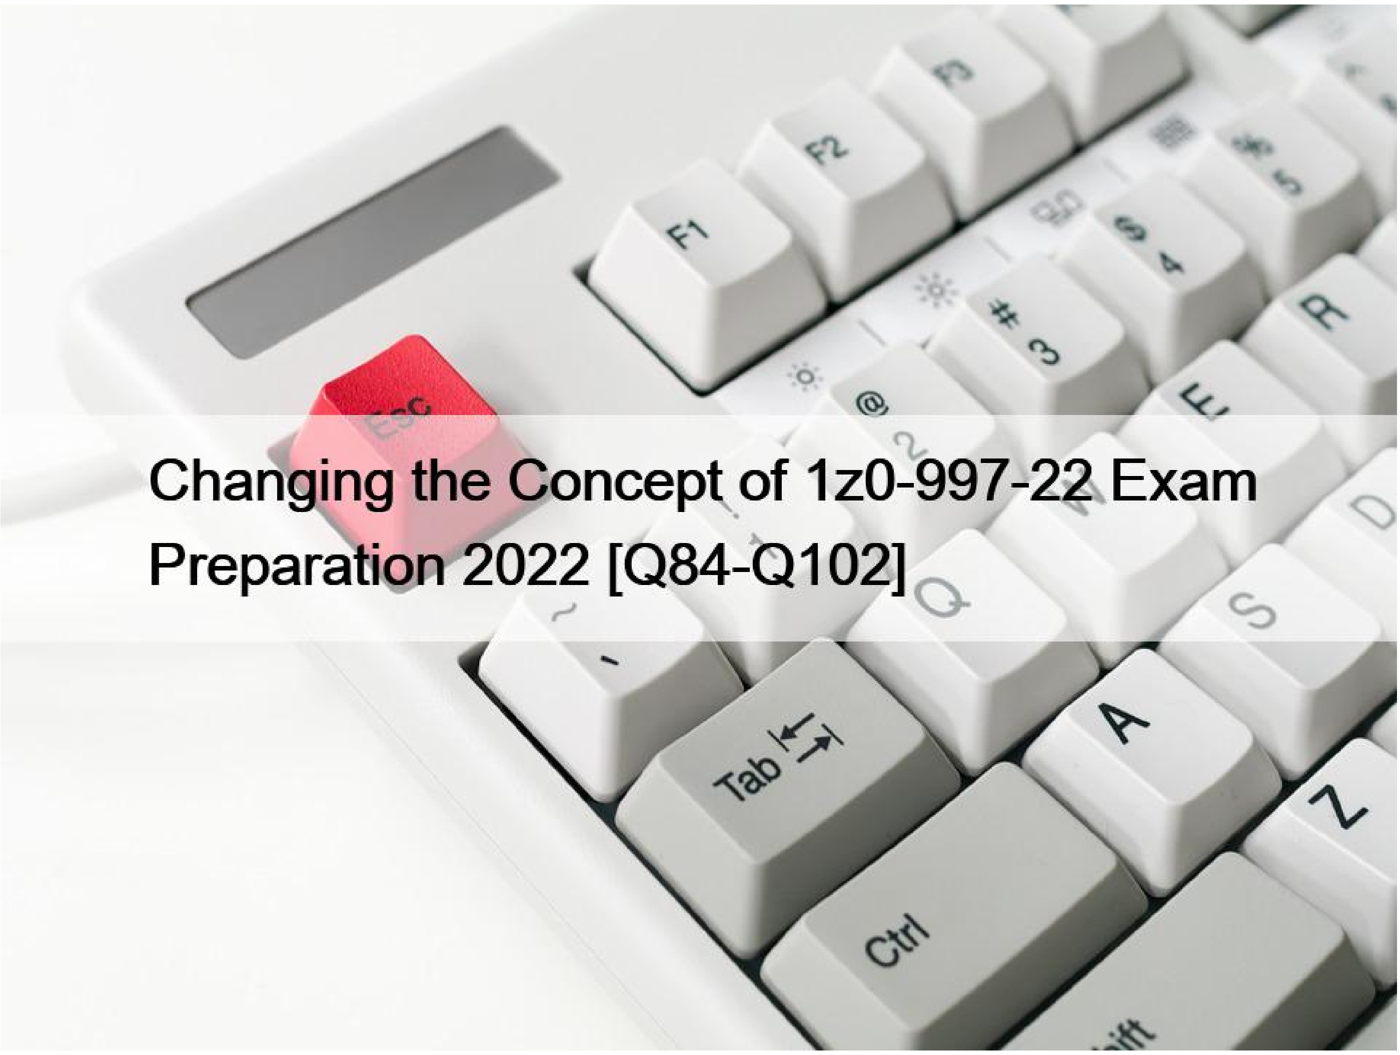

Changing the Concept of 1z0-997-22 Exam Preparation 2022
Getting 1z0-997-22 Certification Made Easy! Get professional help from our 1z0-997-22 Dumps PDF

Oracle 1z0-997-22 Exam Syllabus Topics:

TopicDetailsTopic 1- Manage infrastructure using IaC, OCI CLI, APIs and SDKs- Design and implement hybrid network architectures to meet high availability, bandwidth and latency requirementsTopic 2- Design, implement and operate databases in OCI- Implement and operate solutions in OCI- Evaluate multi-cloud solution architecturesTopic 3- Design strategy for migrating on-premises workloads to OCI- Plan and design solutions to meet business and technical requirements

NO.84 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?
* 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.
* 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.
* 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.
* 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.

NO.85 A large financial company has a web application hosted in their on-premises data center. They are migrating their application to Oracle Cloud Infrastructure (OCI) and require no downtime while the migration is on-going. In order to achieve this, they have decided to divert only 30% of the application works fine, they divert all traffic to OCI.

As a solution architect working with this customer, which suggestion should you provide them?
* Use OCI Traffic management with failover steering policy and distribute the traffic between OC1 and on premises infrastructure.
* Use OCI Traffic management with Load Balancing steering policy and distribute the traffic between OCI and on premises infrastructure.
* Use an OCI load Balancer and distribute the traffic between OCI and on premises infrastructure.
* Use VPN connectivity between on premises Infrastructure and OCI, and create routing tables to distribute the traffic between them.
Traffic Management Steering Policies can account for health of answers to provide failover capabilities, provide the ability to load balance traffic across multiple resources, and account for the location where the query was initiated to provide a simple, flexible and powerful mechanism to efficiently steer DNS traffic.

NO.86 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?
* Create a vault, import your master encryption key into the vault, generate data encryption key, assign data encryption key to the block volume
* 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
* Create a vault, create a master encryption key in the vault, assign this master encryption key to the block volume
* 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
Explanation

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&#8217;s effectively equivalent to a virtual independent HSM.

NO.87 You are working as a cloud consultant for a major media company. In the US and your client requested to consolidate all of their log streams, access logs, application logs, and security logs into a single system.

The client wants to analyze all of their logs In real-time based on heuristics and the result should be validated as well. This validation process requires going back to data samples extracted from the last 8 hours.

What approach should you take for this scenario?
* Create an auto scaling pool of syslog-enabled servers using compute instances which will store the logs In Object storage, then use map reduce jobs to extract logs from Object storage, and apply heuristics on the logs.
* Create a bare-metal instance big enough to host a syslog enabled server to process the logs and store logs on the locally attached NVMe SSDs for rapid retrieval of logs when needed.
* Set up an OCI Audit service and ingest all the API arils from Audit service pragmatically to a client side application to apply heuristics and save the result in an OCI Object storage.
* Stream all the logs and cloud events of Events service to Oracle Streaming Service. Build a client process that will apply heuristics on the logs and store them in an Object Storage.
The Oracle Cloud Infrastructure Streaming service provides a fully managed, scalable, and durable storage solution for ingesting continuous, high-volume streams of data that you can consume and process in real time. Streaming can be used for messaging, ingesting high-volume data such as application logs, operational telemetry, web click-stream data, or other use cases in which data is produced and processed continually and sequentially in a publish-subscribe messaging model.

Streaming Usage Scenarios

Here are some of the many possible uses for Streaming:

Metric and log ingestion: Use the Streaming service as an alternative for traditional file-scraping approaches to help make critical operational data more quickly available for indexing, analysis, and visualization.

Messaging: Use Streaming to decouple components of large systems. Streaming provides a pull/bufferbased communication model with sufficient capacity to flatten load spikes and the ability to feed multiple consumers with the same data independently.
Key-scoped ordering and guaranteed durability provide reliable primitives to implement various messaging patterns, while high throughput potential allows for such a system to scale well.

Web/Mobile activity data ingestion: Use Streaming for capturing activity from websites or mobile apps (such as page views, searches, or other actions users may take). This information can be used for realtime monitoring and analytics, as well as in data warehousing systems for offline processing and reporting.

Infrastructure and apps event processing: Use Streaming as a unified entry point for cloud components to report their life cycle events for audit, accounting, and related activities.

NO.88 You developed a microservices based application that runs on Oracle Cloud Infrastructure (OCI) Container Engine for Kubernetes (OKE). Your security team wants to use SSL termination for this application. What should you do to create a secure SSL

termination for this application using fewest steps?
* Create a self-signed certificate and it&#8217;s corresponding key. Create a Kubernetes secret using the certificate and the key. Then add these an notations to the Kubernetes service:

annotations:

service.beta.kubernetes.io/oci-load-balancer-ssl-ports: &#8220;443&#8221;

service.beta.kubernetes.io/oci-load-balancer-security-list-management-mode:&#8221;Frontend&#8221;
* Generate a self-signed certificate using Let&#8217;s Encrypt. Use that certificate on OCI Load Balancer. Create the Kubernetes service usingthis load balancer.
* Add these annotationsto the Kubernetes service:

annotations:

service.beta.kubernetes.io/oci-load-balancer-ssl-ports: &#8220;443&#8221;

service.beta.kubernetes.io/oci-load-balancer-ssl-secret-key: ssl-secret-key
* Create a self-signed certificate and it&#8217;s corresponding key. Create a Kubernetes secret using then add these annotationsto the Kubernetes service.

Service.beta.kubernete.io/oci-load-balancer-ssl-ports: &#8220;443&#8221;

Service.beta.kubernete.io/oci-load-balancer-tls-secret:SSL-CERTIFICATE-SECRET

NO.89 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?
* 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
* It is possible to spill this fileInto multiple parts using the APIs provided by Object Storage.
* It is possible to split this file into multiple parts using rclone tool provided by Object Storage.
* After initiating a multipart upload by making a CreateMultlPartUpload RESI API Call, the upload remains active until you explicitly commit it or abort.
* Contiguous numbers need to be assigned for each part so that Object Storage constructs the object by ordering, part numbers in ascending order
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.

NO.90 A retail company has several on-premises data centers which span multiple geographical locations. They plan to move some of their applications from on-premises data centers to Oracle Cloud Infrastructure (OCI). For these applications running in OCI, they still need to interact with applications running on their on-premises data centers to Oracle Cloud Infrastructure (OCI). for these applications running in OCI. they still need to interact with applications running on their on-premises data centers. These applications require highly available, fault-tolerant network connections between on premises data centers and OCI.

Which option should you recommend to provide the highest level of redundancy?

* Oracle cloud Infrastructure provides network redundancy by default so that no other operations are required
* If your data centers span multiple, geographical locations, use only the specific IP address as a static route for the specific geographical location
* Set up both IPSec VPN and FastConnect to connect your on premises data centers to Oracle Cloud Infrastructure.
* Use FastConnect private peering only to ensure secure access from your data center to Oracle Cloud Infrastructure
* Set up a single IPSec VPN connection (rom your data center to Oracle Cloud Infrastructure since It is cost effective
If your data centers span multiple geographical locations, we recommend using a broad CIDR (0.0.0.0/0) as a static route in addition to the CIDR of the specific geographical location. This broad CIDR provides high availability and flexibility to your network design.
For instance, the following diagram shows two networks in separate geographical areas that each connect to Oracle Cloud Infrastructure. Each area has a single on-premises router, so two IPSec VPN connections can be created. Note that each IPSec VPN connection has two static routes: one for the CIDR of the particular geographical area, and a broad 0.0.0.0/0 static route.

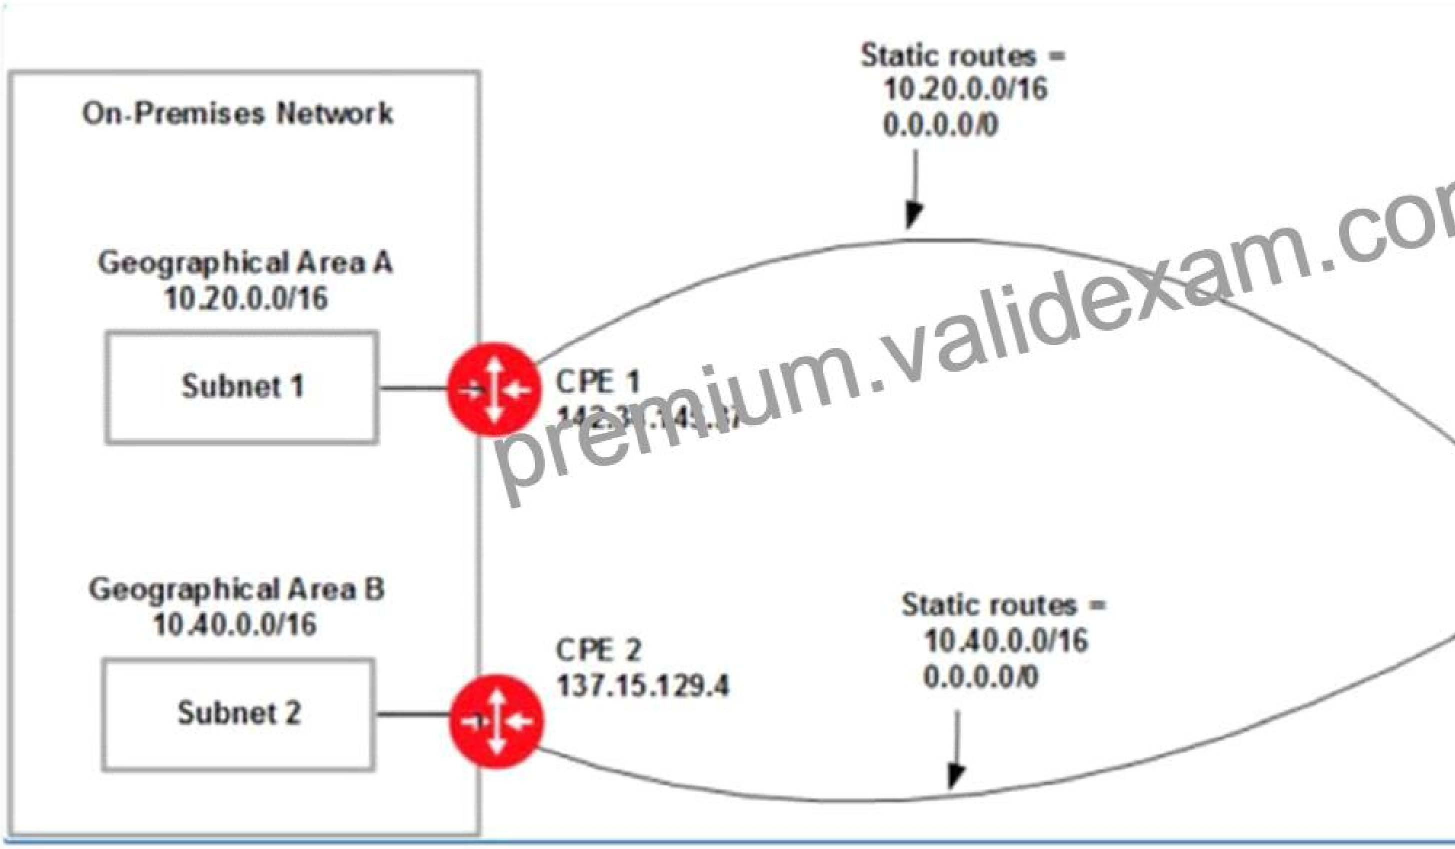

NO.91 An insurance company is storing critical financial data in the Oracle Cloud Infrastructure block volume. This volume is currently encrypted using oracle managed keys. Due to regulatory compliance, the customer wants to encrypt the data using the keys that they can control and not the keys which are controlled by Oracle.

What of the following series of tasks are required to encrypt the block volume using customer managed keys?
* 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.
* Create a vault import your master encryption key into the vault, generate data encryption key, assign data encryption key to the block volume.
* 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.
* Create a vault, create a master encryption key in the vault, assign this master encryption key to the block volume.

NO.92 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?
* Configure hourly block volumes backups using the OCI Command Line Interface (CLI).
* Create a user defined backup policy with a schedule of generating daily backups for block volumes.
* Configure hourly block volumes backups through the OCI Storage Gateway service.
* Create a user defined backup policy with a schedule of generating hourly backups for block volumes.

NO.93 You are currently working for a public health care company based in the United Stats. Their existing patient records runs in an on-premises data center and the customer is sending tape backups offsite as part of their recovery planning.

You have developed an alternative archival solution using Oracle Cloud Infrastructure (OCI) that will save the company a significant amount of mom on a yearly basis. The solution involves storing data in an OCI Object Storage bucket After reviewing your solution with the customer global Compliance (GRC) team they have highlighted the following security requirements:

* All data less than 1 year old must be accessible within 2 hour.

* All data must be retained for at least 10 years and be accessible within 48 hours

* AH data must be encrypted at rest

* No data may be transmitted across the public Internet

Which two options meet the requirements outlined by the customer GRC team?
* Provision a FastConnect link to the closest OCI region and configure a private peering virtual circuit.
* Create an OCI Object Storage Standard tier bucket Configure a lifecycle policy to archive any object that Is older than 365 days
* Create a VPN connection between your on premises data center and OCI. Create a Virtual Cloud Network (VCN) along with an OCI Service Gateway for OCI Object Storage.
* Provision a FastConnect link to the closest OCI region and configure a public peering virtual circuit
* Create an OCI Object Storage Standard tier bucket. Configure a lifecycle policy to delete any object that is older than 7 years
The Oracle Services Network is a conceptual network in Oracle Cloud Infrastructure that is reserved for Oracle services. These services have public IP addresses that you typically reach over the internet. However, you can access the Oracle Services Network without the traffic going over the internet. There are different ways, depending on which of your hosts need the access:

Hosts in your on-premises network:

&#8211; Private access through a VCN with FastConnect private peering or VPN Connect: The on-premises hosts use private IP addresses and reach the Oracle Services Network by way of the VCN and the VCN&#8217;s service gateway.

&#8211; Public access with FastConnect public peering: The on-premises hosts use public IP addresses.

regarding which Fastconnect Public peering: To access public services in Oracle Cloud Infrastructure without using the internet. For example, Object Storage, the Oracle Cloud Infrastructure Console and APIs, or public load balancers in your VCN. Communication across the connection is with IPv4 public IP addresses. Without FastConnect, the traffic destined for public IP addresses would be routed over the internet. With FastConnect, that traffic goes over your private physical connection.

so Answer 4 will be the best answer that meets the customer requirement A service gateway lets your virtual cloud network (VCN) privately access specific Oracle services without exposing the data to the public internet. No internet gateway or NAT is required to reach those specific services. The resources in the VCN can be in a private subnet and use only private IP addresses. The traffic from the VCN to the Oracle service travels over the Oracle network fabric and never traverses the internet.

Object Lifecycle Management lets you automatically manage the archiving and deletion of objects. By using Object Lifecycle Management to manage your Object Storage and Archive Storage data, you can reduce your storage costs and the amount of time you spend managing data.

NO.94 A manufacturing company is planning to migrate their on-premises database to OCI and has hired you for the migration. Customer has provided following information regarding their existing onpremises database:

Database version, host operating system and version, database character set, storage for data staging, acceptable length of system outage.

What additional information do you need from customer in order to recommend a suitable migration method? Choose two
* Elapsed time since database was last patched
* On-premises host operating system and version
* Number of active connections
* Data types used in the on-premises database
* Top 5 longest running queries
Not all migration methods apply to all migration scenarios. Many of the migration methods apply only if specific characteristics of the source and destination databases match or are compatible. Moreover, additional factors can affect which method you choose for your migration from among the methods that are technically applicable to your migration scenario.

Some of the characteristics and factors to consider when choosing a migration method are:

On-premises database version

Database service database version

On-premises host operating system and version

On-premises database character set

Quantity of data, including indexes

Data types used in the on-premises database

Storage for data staging

Acceptable length of system outage

Network bandwidth

NO.95 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?
* Create a default tag for each compartment, which ensure that appropriate tags are applied at the time of resource creation.
* Create an OCI Identity and Access Management policy requiring users to tag resources with their user name.
* Create an OCI Identity and Access Management policy to automatically tag a resource with the user name.
* Create tag variables to automatically tag a resource with the user name.

NO.96 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 &#8216;Volume Groups&#8217; 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?
* The Identity and Access Management policy allowing backup failed to move when the compartment was moved.
* You are exceeding your volume group backup quota configured.
* 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.
* A compute instance with multiple block volumes attached cannot move when a compartment is moved.

NO.97 You are responsible for migrating your on-premises legacy databases on 11.2.0.4 version to Autonomous Transaction Processing &#8211; Dedicated (ATP-D) in Oracle Cloud Infrastructure (OCI). As a solution architect, you need to plan your migration approach.

Which three options do you need to implement together to migrate your on-premises databases to OCI?
* Retain all legacy structures and unsupported features (e.g. legacy LOBs) in the on-premises databases for migration.
* Use Oracle Data Guard to keep on-premises database always active during migration.
* Launch Autonomous Transaction Processing &#8211; Dedicated database in OCI.
* Retain changes to Oracle shipped privileges, stored procedures or views in the on-premises databases.
* Convert on-premises databases to PDB, upgrade to 19c, and encrypt.
* Use Oracle GoldenGate replication to keep on-premises database online during migration.

NO.98 You are working as a cloud engineer for an IoT startup company which is developing a health monitoring pet collar for dogs and cats. The company collects biometric Information of the pet every second and then sends it to Oracle Cloud Infrastructure (OCI) Your task is to come up with an architecture which will accept and process the monitoring data as well as provide complete trends and health reports to the pet owners. The portal should be highly available, durable, and scalable with an additional feature for showing real time biometric data analytics.

which architecture will help you meet this requirement?
* Use OCI Streaming Service to collect the incoming biometric data. Use Oracle Functions to process the date and show the results on a real-time dashboard and store the results lo OCI Object Storage Store the data In OCI Autonomous Data warehouse (ADW) to handle analytics.
* Launch an open source Hadoop cluster to collect the Incoming biometrics data Use an Open source Fluentd cluster to analyze the-data me results to OCI Autonomous Transaction Processing (ADW)to handle complex analytics
* Create an OCI Object Storage bucket to collect the incoming biometric data from the smart pet collar Fetch the data horn OC Object storage to OCI Autonomous Data Warehouse (ADW) every day and run analytics Jobs with it
* Use OCI Streaming Service to collect the incoming biometric data. Use an open source Hadoop cluster to analyze the data horn streaming service. Store the results to OCI Autonomous Data warehouse (ADW) to handle complex analytics.

NO.99 Which of the following is NOT a good use case for using the functionality available in the Oracle Cloud Infrastructure (OCI) Events service?

* Publish all events in a specific compartment to Oracle Streaming service for later analysis.
* Triggers Function using Oracle Functions when new files are uploaded in an OCI Object Storage bucket.
* Publish a notification when long lived tasks complete, such as OCI Autonomous Database backup completion.
* Capture Monitoring Alarms and invoke Autoscaling of compute instances.
* Trigger a notification when a function completes its execution.

NO.100 You are the Solution Architect that designed this Oracle Cloud Infrastructure (OCI) compartment layout for your organization:

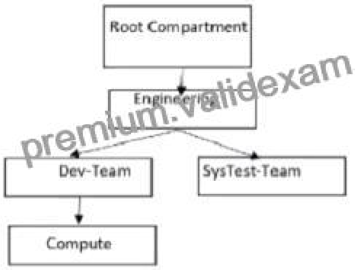

The development team has deployed quite a few instances under &#8216;Compute&#8217; Compartment and the operations team needs to list the Instances under the same compartment for their testing. Both teams, development and operations are part of a group called &#8216;Eng-group&#8217; You have been looking for an option to allow the operations team to list the instances without access any confidential information or metadata of resources.

Which IAM policy should you write based on these requirements?

* Allow group Eng-group to inspect instance-family in compartment Dev-Team:Compute and attach the policy to &#8216;Engineering&#8217; Compartment
* Allow group Eng-group to inspect instance-family in compartment Dev-Team: Compute and attach the policy to &#8216;SysTest Team&#8217; Compartment
* Allow group Eng-group to read instance-family in compartment Compute and attach the policy to &#8216;Engineering&#8217; Compartment.
* Allow group Eng-group to read instance-family in compartment Dev-Team-.Compute and attach the policy to&#8217;Dev-Team&#8217;
Policy Attachment

When you create a policy you must attach it to a compartment (or the tenancy, which is the root compartment). Where you attach it controls who can then modify it or delete it. If you attach it to the tenancy (in other words, if the policy is in the root compartment), then anyone with access to manage policies in the tenancy can then change or delete it. Typically that&#8217;s the Administrators group or any similar group you create and give broad access to. Anyone with access only to a child compartment cannot modify or delete that policy.

When you attach a policy to a compartment, you must be in that compartment and you must indicate directly in the statement which compartment it applies to. If you are not in the compartment, you&#8217;ll get an error if you try to attach the policy to a different compartment. Notice that attachment occurs during policy creation, which means a policy can be attached to only one compartment.

Policies and Compartment Hierarchies

a policy statement must specify the compartment for which access is being granted (or the tenancy).

Where you create the policy determines who can update the policy. If you attach the policy to the compartment or its parent, you can simply specify the compartment name. If you attach the policy further up the hierarchy, you must specify the path. The format of the path is each compartment name (or OCID) in the path, separated by a colon:

<compartment_level_1>:<compartment_level_2>: . . . <compartment_level_n> to allow action to compartment Compute so you need to set the compartment PATH as per where you attach the policy as below examples if you attach it to Root compartment you need to specify the PATH as following Engineering:Dev-Team:Compute if you attach it to Engineering compartment you need to specify the PATH as following Dev-Team:Compute if you attach it to Dev-Team or Compute compartment you need to specify the PATH as following Compute Note : in the Policy inspect verb that give the Ability to list resources, without access to any confidential information or user-specified metadata that may be part of that resource.

NO.101 An Oracle Cloud Infrastructure (OCI) Public Load Balancer&#8217;s SSL certificate is expiring soon. You noticed the Load Balancer is configured with SSL Termination only. When the certificate expires, data traffic can be interrupted and security compromised.

What steps do you need to take to prevent this situation?
* Add the new SSL certificate to the Load Balancer, update backend servers to work with a new certificate and edit listeners so they can use the new certificate bundle.
* Add the new SSL certificate to the Load Balancer, update listeners and backend sets so they can use the new certificate bundle.
* Add the new SSL certificate to the Load Balancer and implement end to end SSL so it can encrypt the traffic from clients all the way to the backend servers.
* Add the new SSL certificate to the Load Balancer and update backend servers to use the new certificate bundle.
* Add the new SSL certificate to the Load Balancer and update listeners to use the new certificate bundle.
https://docs.cloud.oracle.com/en-us/iaas/Content/Balance/Tasks/managingcertificates.htm

NO.102 A retail company has recently adopted a hybrid architecture. They have the following requirements for their end-to-end Connectivity model between their on-premises data center and Oracle Cloud Infrastructure (OC1) region

* Highly available connection with service level redundancy

* Dedicated network bandwidth with low latency

Which connectivity setup is the most cost effective solution for this scenario?
* Setup IPsec VPN as your primary connection, and a FastConnect virtual circuit as a backup connection. Use separate edge devices in your on-premises data canter for each connection from your edge devices, advertise more specific routes IPSec VPN, and specific routes through the backup FastConnect virtual circuit.
* Setup FastConnect virtual circuit as your primary connection, and a second FastConnect virtual circuit as a backup connection. Use separate edge devices in your FastConnect physical connectivity is redundant Use a single edge device in your on premises data center for each connection From yc device, advertise more specific routes via primary FastConnect virtual circuit, and less specific routes through t backup FastConnect circuit.
* Setup FastConnect virtual circuit as your primary connection, and an IPSec VPN as a backup connection. Use separate edge devices in your on-premises data center for each connection. From your edge devices, advertise more specific routes through FastConnect virtual circuit, and more specific routes through the backup IPSec VPN path.
* Setup IPSec VPN as your primary connection, and a second IPSec VPN as a backup connection. Use separate edge devices in your on p data center for each connection. From your edge devices, advertise more specific routes via primary IPSec VPN. and less specific rod the backup IPSec VPN.
there are two main requirements for this Customer

First Highly available connection with service level redundancy and that can achieve by

1- VPN Connect with a Redundant Customer Edge Device

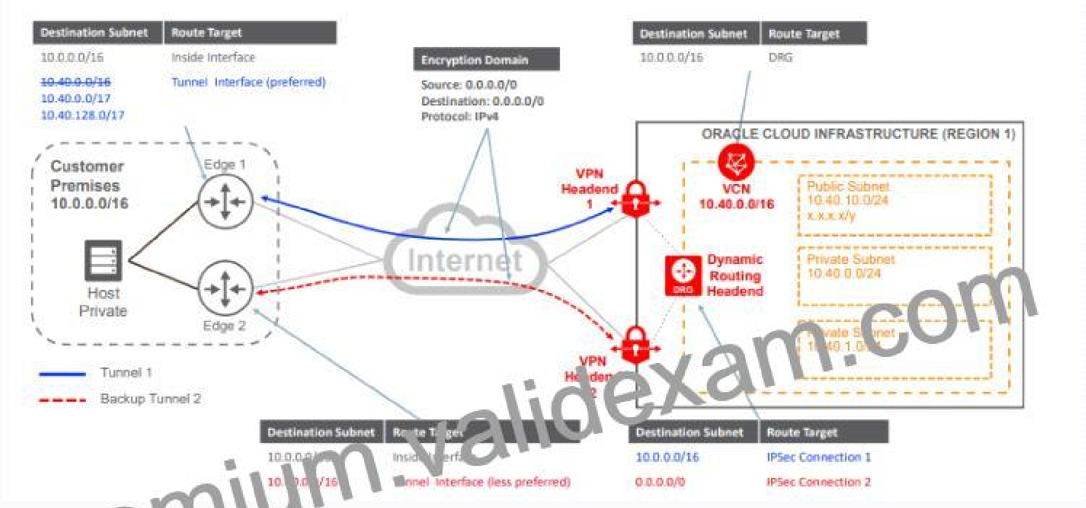
Customer Premises 10.0.0.0/16
Edge 1
Edge 2
Host Private
Internet
VPN Headend 1
VPN Headend 2
ORACLE CLOUD INFRASTRUCTURE (REGION 1)
VCN 10.40.0.0/16
Dynamic Routing Headend
DRG
Public Subnet 10.40.10.0/24 x.x.x.x/y
Private Subnet 10.40.0.0/24
Private Subnet 10.40.1.0/24
Tunnel 1
Backup Tunnel 2

Destination Subnet	Route Target
10.0.0.0/16	Inside Interface
~~10.40.0.0/16~~ 10.40.0.0/17 10.40.128.0/17	Tunnel Interface (preferred)

Encryption Domain
Source: 0.0.0.0/0 Destination: 0.0.0.0/0 Protocol: IPv4

Destination Subnet	Route Target
10.0.0.0/16	DRG

Destination Subnet	Route Target
10.0.0.0/16	Inside Interface
10.40.0.0/16	Tunnel Interface (less preferred)

Destination Subnet	Route Target
10.0.0.0/16	IPSec Connection 1
0.0.0.0/0	IPSec Connection 2

2- FastConnect Plus a Single VPN Connect Connection

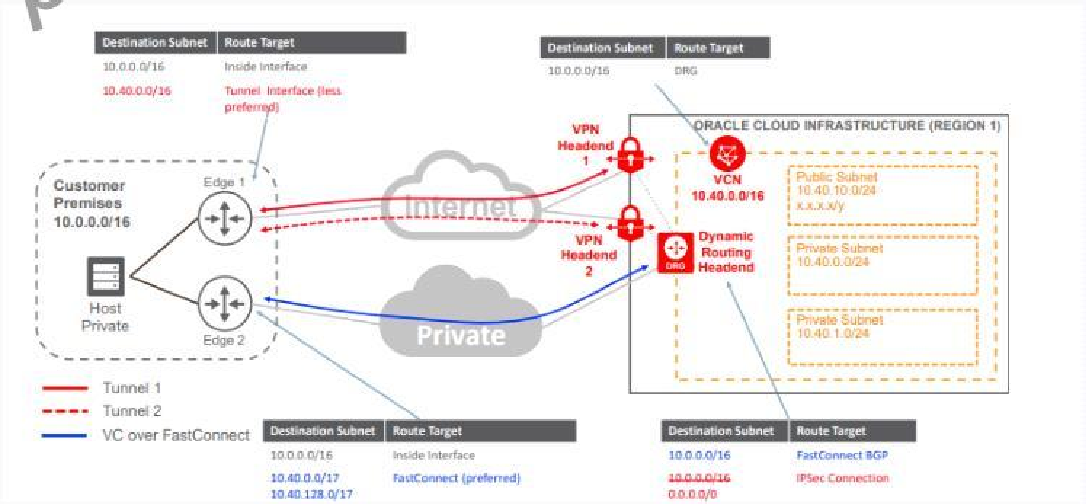
Customer Premises 10.0.0.0/16
Edge 1
Edge 2
Host Private
Internet
Private
VPN Headend 1
VPN Headend 2
ORACLE CLOUD INFRASTRUCTURE (REGION 1)
VCN 10.40.0.0/16
Dynamic Routing Headend
DRG
Public Subnet 10.40.10.0/24 x.x.x.x/y
Private Subnet 10.40.0.0/24
Private Subnet 10.40.1.0/24
Tunnel 1
Tunnel 2
VC over FastConnect

Destination Subnet	Route Target
10.0.0.0/16	Inside Interface
10.40.0.0/16	Tunnel Interface (less preferred)

Destination Subnet	Route Target
10.0.0.0/16	DRG

Destination Subnet	Route Target
10.0.0.0/16	Inside Interface
10.40.0.0/17 10.40.128.0/17	FastConnect (preferred)

Destination Subnet	Route Target
10.0.0.0/16	FastConnect BGP
~~10.0.0.0/16~~ 0.0.0.0/0	IPSec Connection

3- Redundant FastConnect

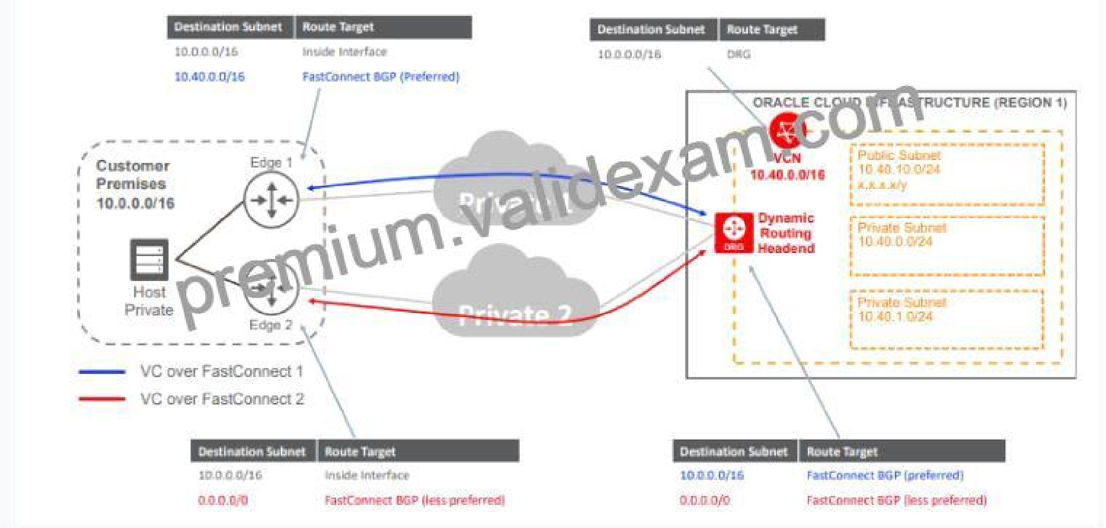
Customer Premises 10.0.0.0/16
Edge 1
Edge 2
Host Private
Private
Private 2
ORACLE CLOUD INFRASTRUCTURE (REGION 1)
VCN 10.40.0.0/16
Dynamic Routing Headend
DRG
Public Subnet 10.40.10.0/24 x.x.x.x/y
Private Subnet 10.40.0.0/24
Private Subnet 10.40.1.0/24
VC over FastConnect 1
VC over FastConnect 2

Destination Subnet	Route Target
10.0.0.0/16	Inside Interface
10.40.0.0/16	FastConnect BGP (Preferred)

Destination Subnet	Route Target
10.0.0.0/16	DRG

Destination Subnet	Route Target
10.0.0.0/16	Inside Interface
0.0.0.0/0	FastConnect BGP (less preferred)

Destination Subnet	Route Target
10.0.0.0/16	FastConnect BGP (preferred)
0.0.0.0/0	FastConnect BGP (less preferred)

1z0-997-22 Exam Crack Test Engine Dumps Training With 145 Questions:
https://www.validexam.com/1z0-997-22-latest-dumps.html]