

1z0-997-20 Dumps

Oracle Cloud Infrastructure 2020 Architect Professional

<https://www.certleader.com/1z0-997-20-dumps.html>



NEW QUESTION 1

A data analytics company has been building its now generation big data and analytics platform on Oracle Cloud Infrastructure (OCI). They need a storage service that provide the scale and performance that their big data applications require such as high throughput to compute nodes with low latency file operations in addition, their data needs to be stored redundantly across multiple nodes in a single availability domain and allows concurrent connections from multiple compute instances hosted on multiple availability domains.

Which OCI storage service can you use to meet his requirement?

- A. Object Storage
- B. File System Storage
- C. Archive storage
- D. Block Volume

Answer: B

Explanation:

Oracle Cloud Infrastructure File Storage service provides a durable, scalable, secure, enterprise-grade network file system. You can connect to a File Storage service file system from any bare metal, virtual machine, or container instance in your Virtual Cloud Network (VCN). You can also access a file system from outside the VCN using Oracle Cloud Infrastructure FastConnect and Internet Protocol security (IPSec) virtual private network (VPN).

Use the File Storage service when your application or workload includes big data and analytics, media processing, or content management, and you require Portable Operating System Interface (POSIX)- compliant file system access semantics and concurrently accessible storage. The File Storage service is designed to meet the needs of applications and users that need an enterprise file system across a wide range of use cases

NEW QUESTION 2

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?

- A. Use OCI Streaming Service to collect the incoming biometric data
- B. Use Oracle Functions to process the data and show the results on a real-time dashboard and store the results to OCI Object Storage. Store the data in OCI Autonomous Data Warehouse (ADW) to handle analytics.
- C. Launch an open source Hadoop cluster to collect the incoming biometrics data. Use an open source Fluentd cluster to analyze the data. Send results to OCI Autonomous Transaction Processing (ATP) to handle complex analytics
- D. Create an OCI Object Storage bucket to collect the incoming biometric data from the smart pet collar. Fetch the data from OCI Object Storage to OCI Autonomous Data Warehouse (ADW) every day and run analytics jobs with it
- E. Use OCI Streaming Service to collect the incoming biometric data
- F. Use an open source Hadoop cluster to analyze the data from streaming service
- G. Store the results to OCI Autonomous Data Warehouse (ADW) to handle complex analytics.

Answer: A

NEW QUESTION 3

Your organization is planning on using Oracle Cloud Infrastructure (OCI) File Storage Service (FSS). You will be deploying multiple compute instances in Oracle Cloud Infrastructure (OCI) and mounting the file system to these compute instances.

The file system will hold payment data processed by a Database instance and utilized by compute instances to create an overall inventory report. You need to restrict access to this data for specific compute instances and must be allowed/blocked per compute instance's CIDR block.

Which option can you use to secure access?

- A. Create a new VCN security list, choose SOURCE TYPE as Service and SOURCE SERVICE as FS
- B. Add stateless ingress and egress rules for specific IP address and CIDR blocks.
- C. Use 'Export option' feature of FSS to restrict access to the mounted file systems.
- D. Create and configure OCI Web Application Firewall service with built-in DNS-based intelligent routing.
- E. Use stateless Security List rule to restrict access from known IP addresses only.

Answer: B

NEW QUESTION 4

All three Data Guard configurations 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 RPO 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 ASYNC mode between two regions.

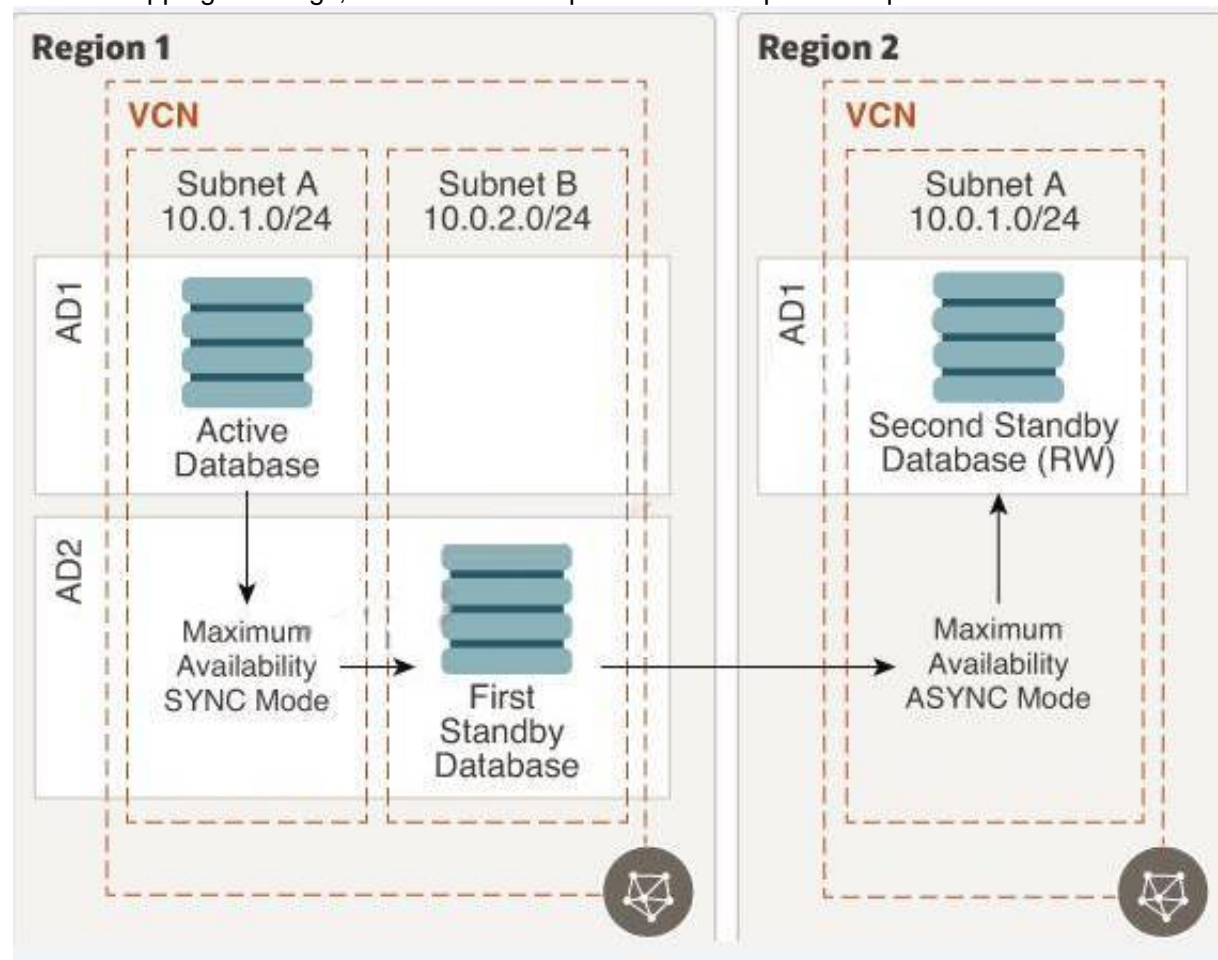
Answer: D

Explanation:

<https://docs.cloud.oracle.com/en-us/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

NEW QUESTION 5

You are working with a customer who needs to attach an Oracle Cloud Infrastructure (OCI) block volume to a VM instance with read/write access type. The customer wants to know if the number of IOPS and throughput performance differs between the following two choices:

- Option A: attach a single 1 TB block volume to the VM instance
 - Option B: attach two separate 500 GB block volumes in a RAID 0 array configuration to the VM instance
- You can assume that the customer is using iSCSI attachment type to attach the volumes to the instance. In addition, you can assume 1 MB block size for throughput and 4 KB block size for IOPS consideration. How should you respond to the customer?

- A. Option B provides higher level of throughput, but lower level of IOPS performance.
- B. Both options provide the same number of IOPS and throughput performance.
- C. Option A provides better IOPS, but lower throughput performance.
- D. Option B provides better IOPS and throughput performance.

Answer: B

NEW QUESTION 6

Your company will soon start moving critical systems into Oracle Cloud Infrastructure (OCI) platform. These systems will reside in the us-phoenix-1 and us-ashburn-1 regions. As part of the migration planning, you are reviewing the company's existing security policies and written guidelines for the OCI platform usage within the company. You have to work with the company managed key.

Which two options ensure compliance with this policy?

- A. When you create a new compute instance through OCI console, you use the default options for "configure boot volume" to speed up the process to create this compute instance.
- B. When you create a new block volume through OCI console, select Encrypt using Key Management checkbox and use encryption keys generated and stored in OCI Key Management Service.
- C. When you create a new compute instance through OCI console, you use the default shape to speed up the process to create this compute instance.
- D. When you create a new OCI Object Storage bucket through OCI console, you need to choose "ENCRYPT USING CUSTOMER-MANAGED KEYS" option.
- E. You do not need to perform any additional actions because the OCI Block Volume service always encrypts all block volumes, boot volumes, and volume backups at rest by using the Advanced Encryption Standard (AES) algorithm with 256-bit encryption.

Answer: BD

Explanation:

Block Volume Encryption

By default all volumes and their backups are encrypted using the Oracle-provided encryption keys. Each time a volume is cloned or restored from a backup the volume is assigned a new unique encryption key.

You have the option to encrypt all of your volumes and their backups using the keys that you own and manage using the Vault service. If you do not configure a volume to use the Vault service or you later unassign a key from the volume, the Block Volume service uses the Oracle-provided encryption key instead.

Create Block Volume

Size: 50 GB
Size must be between 10 GB and 32,768 GB (32 TB). Volume performance varies with volume size.

COMPARTMENT FOR BACKUP POLICIES: techoci (root)

BACKUP POLICY: Select a Backup Policy

VOLUME PERFORMANCE: Balanced
Balanced choice for most workloads including those that perform random I/O such as boot disks. [Learn more](#)
IOPS: 3000 IOPS (50 IOPS/GB)
Throughput: 24 MB/s (400 KB/s/GB)

ENCRYPTION: ☒ ENCRYPT USING CUSTOMER MANAGED KEYS
Leaves all encryption-related matters to Oracle.
☐ ENCRYPT USING ORACLE MANAGED KEYS
Requires you to have access to a valid Key Management key.

VAULT COMPARTMENT: techoci (root) VAULT: demo_vault

MASTER ENCRYPTION KEY COMPARTMENT: techoci (root) MASTER ENCRYPTION KEY: Demo_Key

This applies to both encryption at-rest and in-transit encryption. Object Storage Encryption

Object Storage employs 256-bit Advanced Encryption Standard (AES-256) to encrypt object data on the server. Each object is encrypted with its own data encryption key. Data encryption keys are always encrypted with a master encryption key that is assigned to the bucket. Encryption is enabled by default and cannot be turned off. By default, Oracle manages the master encryption key. However, you can optionally configure a bucket so that it's assigned an Oracle Cloud Infrastructure Vault master encryption key that you control and rotate on your own schedule.

Encryption: Buckets are encrypted with keys managed by Oracle by default, but you can optionally encrypt the data in this bucket using your own Vault encryption key. To use Vault for your encryption needs, select

Encrypt Using Customer-Managed Keys. Then, select the Vault Compartment and Vault that contain the master encryption key you want to use. Also select the Master Encryption Key Compartment and Master Encryption Key.

Create Bucket

BUCKET NAME: bucket-

STORAGE TIER: STANDARD
Storage tier for a bucket can only be specified during creation. Once set, you cannot change the storage tier in which a bucket resides.

OBJECT EVENTS: ☐ EMIT OBJECT EVENTS

OBJECT VERSIONING: ☐ ENABLE OBJECT VERSIONING

ENCRYPTION: ☒ ENCRYPT USING CUSTOMER MANAGED KEYS
Requires you to have access to a valid Key Management key. [Learn More](#)
☐ ENCRYPT USING ORACLE MANAGED KEYS
Leaves all encryption-related matters to Oracle.

VAULT COMPARTMENT: techoci (root) VAULT: demo_vault

MASTER ENCRYPTION KEY COMPARTMENT: techoci (root) MASTER ENCRYPTION KEY: Demo_Key

NEW QUESTION 7

An organization has its IT infrastructure in a hybrid setup with an on-premises environment and an Oracle Cloud Infrastructure (OCI) Virtual Cloud Network (VCN) in the us-phoenix-1 region. The on-premise applications communications with compute instances inside the VPN over a hardware VPN connection. They are looking to implement an Intrusion detected and Prevention (IDS/IPS) system for their OCI environment. This platform should have the ability to scale to thousands of compute of instances running inside the VCN.

How should they architect their solution on OCI to achieve this goal?

- A. Set up an OCI Private Load Balance! and configure IDS/IPS related health checks at TCP and/or HTTP level to inspect traffic
- B. Configure each host with an agent that collects all network traffic and sends that traffic to the IDS/IPS platform to inspection
- C. There is no need to implement an IPS/IDS system as traffic coming over IPsec VPN tunnels is already encrypt
- D. Configure autoscaling on a compute Instance pool and set vNIC to promiscuous mode to capture traffic across the vcn and send it IDS/IPS platform for inspection.

Answer: B

Explanation:

In Transit routing through a private IP in the VCN you set up an instance in the VCN to act as a firewall or intrusion detection system to filter or inspect the traffic between the on-premises network and Oracle Services Network.

The Networking service lets you implement network security functions such as intrusion detection, application-level firewalls. In fact, the IDS model can be host-based IDS (HIDS) or network-based IDS

(NIDS). HIDS is installed at a host to periodically monitor specific system logs for patterns of intrusions. In

contrast, an NIDS sniffs the traffic to analyze suspicious behaviors. A signature-based NIDS (SNIDS) examines the traffic for patterns of known intrusions. SNIDS can quickly and reliably diagnose the attacking techniques and security holes without generating an over-whelming number of false alarms because SNIDS relies on known signatures.

However, anomaly-based NIDS (ANIDS) detects unusual behaviors based on statistical methods. ANIDS could detect symptoms of attacks without specific knowledge of details. However, if the training data of the normal traffic are inadequate, ANIDS may generate a large number of false alarms.

NEW QUESTION 8

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?

- A. Use OCI Traffic management with failover steering policy and distribute the traffic between OCI and on premises infrastructure.
- B. Use OCI Traffic management with Load Balancing steering policy and distribute the traffic between OCI and on premises infrastructure.
- C. Use an OCI load Balancer and distribute the traffic between OCI and on premises infrastructure.
- D. Use VPN connectivity between on premises Infrastructure and OCI, and create routing tables to distribute the traffic between them.


Answer: B

Explanation:

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.

NEW QUESTION 9

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 'Compute' 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 'Eng-group' 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?

- A. Allow group Eng-group to inspect instance-family in compartment Dev-Team:Compute and attach the policy to 'Engineering' Compartment
- B. Allow group Eng-group to inspect instance-family in compartment Dev-Team: Compute and attach the policy to 'SysTest Team' Compartment
- C. Allow group Eng-group to read instance-family in compartment Compute and attach the policy to 'Engineering' Compartment.
- D. Allow group Eng-group to read instance-family in compartment Dev-Team-.Compute and attach the policy to 'Dev-Team'

Answer: A

Explanation:

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'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'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.

NEW QUESTION 10

A retail company runs their online shopping platform entirely on Oracle cloud Infrastructure (OCI). This is a 3-tier web application that Includes a Mbps Load Balancer. Virtual Machine Instances for web and an Oracle DB Systems Virtual Machine Due to unprecedented growth, they noticed an Increase in the Incoming traffic to their website and all users start getting 503 (Service Unavailable) errors.

What is the potential problem in this scenario?

- A. The Load Balancer health check status Indicates critical situation for half of the backend web servers
- B. All the web servers are too busy and not able to answer any request from users.

- C. The Database Is down hence users can not access the web site
- D. The Traffic Management Policy is not set to load Balancer the traffic to the web servers.
- E. You did not configure a Service Gateway to allow connection between web servers and load Balance

Answer: B

Explanation:

A 503 Service Unavailable Error is an HTTP response status code indicating that a server is temporarily unable to handle the request. This may be due to the server being overloaded or down for maintenance.

NEW QUESTION 10

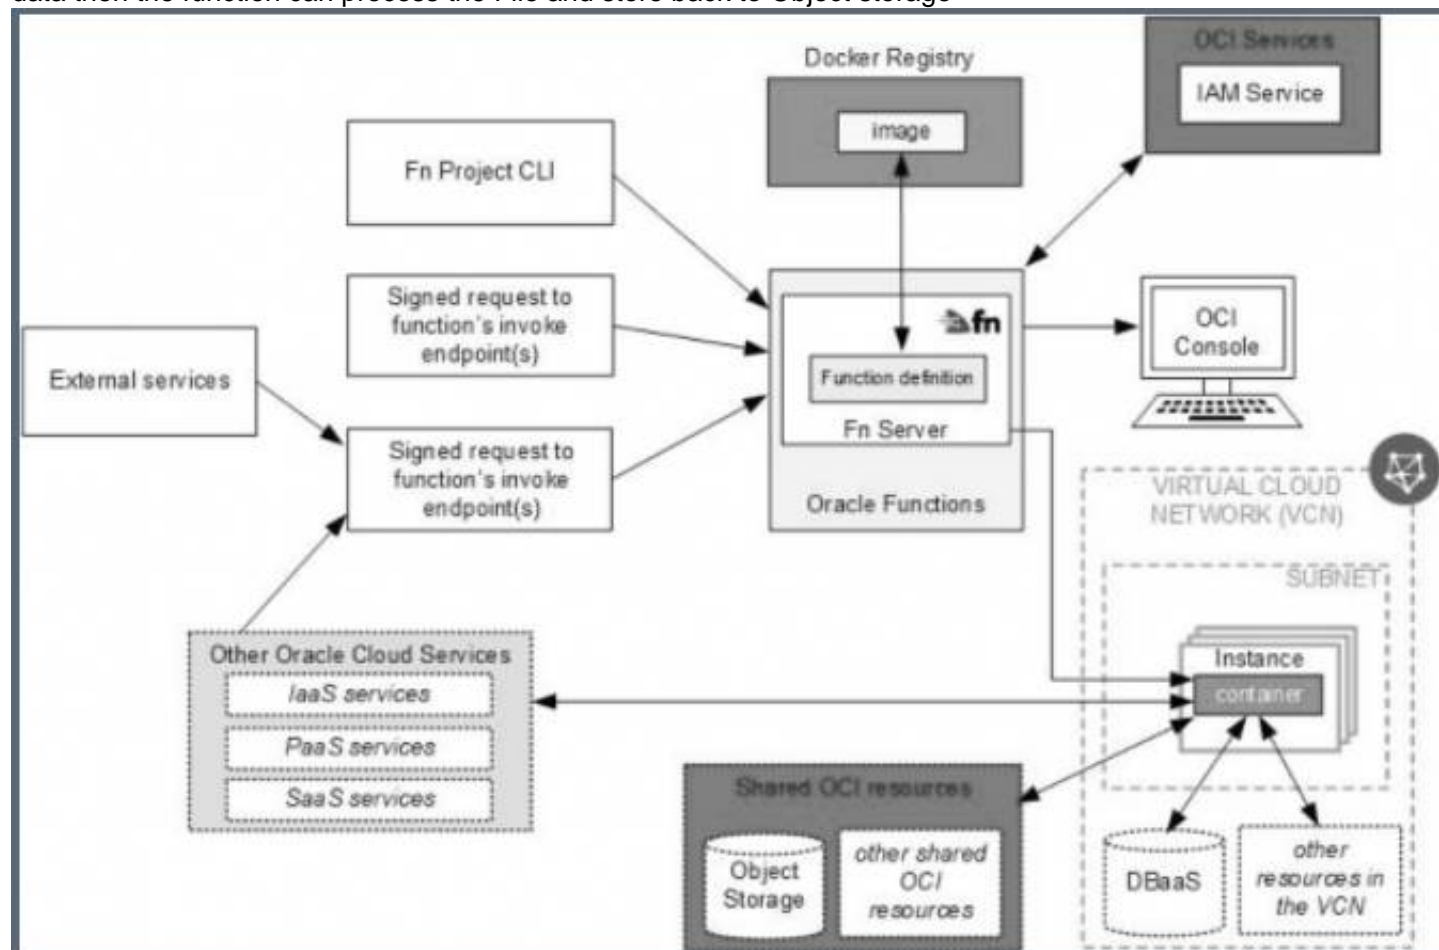
You want to automate the processing of new Image files to generate thumbnails. the expected rate is 10 new files every hour. Which of the following is the most cost effective option to meet this requirement in Oracle Cloud Infrastructure (OCI)?

- A. Upload files to an OCI Object storage bucket
- B. Every time a file is uploaded, an event is emitted
- C. Write a rule to filter these events with an action to trigger a function in Oracle Function
- D. The function processes the image in the file and stores the thumbnails back in an Object storage bucket.
- E. Upload files to an OCI Object storage bucket
- F. Every time a file is uploaded, trigger an event with an action to provision a compute instance with a cloud-init script to access the file, process it and store it back in an Object storage bucket
- G. Terminate the instance using Autoscaling policy after the processing is finished.
- H. Build a web application to ingest the files and save them to a NoSQL Database
- I. Configure OCI Events service to trigger a notification using Oracle Notification Service (ONS). ONS invokes a custom application to process the image files to generate thumbnail
- J. Store thumbnails in a NoSQL Database table.
- K. Upload all files to an Oracle Streaming Service (OSS) stream
- L. Set up a cron job to invoke a function in Oracle Functions to fetch data from the stream
- M. Invoke another function to process the image files and generate thumbnail
- N. Store thumbnails in another OSS stream.

Answer: A

Explanation:

You can invoke a function that you've deployed to Oracle Functions by triggered by an event in the Events service when update the Object storage to fetch the data then the function can process the File and store back to Object storage



NEW QUESTION 14

You are tasked with building a highly available, fault tolerant web application for your current employer. The security team is concerned about an increase in malicious web-based attacks across the internet and asked what you can do to add a higher level of security to the website. How should you architect the solution on Oracle Cloud Infrastructure (OCI) to meet all requirements defined by your organization? (Choose the best answer.)

- A. Deploy at least 3 web application servers, each in a different fault domain, using a regional private subne
- B. Place a public load balancer in a regional public subnet and create a backend set for all of the web application server
- C. Deploy a Web Application Firewall (WAF) and configure the load balancer public IP address as the origin.
- D. Deploy at least 3 web application servers, each in a different fault domain, using a regional private subne
- E. Place a public load balancer in a regional public subnet and create a backend set for all of the web application server
- F. Create a Geolocation steering policy in Traffic Management and add an answer pool that directs to the public IP address of the load balance
- G. Configure a global catch-all rule to use this answer pool.
- H. Deploy at least 3 web application servers, each in a different fault domain, using a regional public subne
- I. Ensure that each web application server is assigned a public IP address
- J. Deploy a Web Application Firewall (WAF) and configure one Origin for each public IP address.

- K. Deploy at least 3 web application servers, each in a different fault domain, using a regional public subne
- L. Use the OCI Traffic Management service to create a load balancing policy that will resolve DNS evenly between all web servers.

Answer: A

NEW QUESTION 16

Which of the following is NOT a good use case for the Oracle Cloud Infrastructure (OCI) Streaming service?

- A. Meeting compliance requirements for data to remain unchanged over a long time, so that it can be retrieved for audit purposes.
- B. Messaging with a pull-based communication model and the ability to feed multiple consumers with the same data independently.
- C. Ingesting metric and log data to help make critical operational data more quickly available for indexing, analysis, and visualization.
- D. Providing a unified entry point for cloud components to report their life cycle events for audit, accounting, and related activities.

Answer: A

NEW QUESTION 17

You work for a large bank where your main application is a payment processing gateway API. You deployed the application on Oracle Container Engine for Kubernetes (OKE) and used API Gateway with several policies to control the access of the API endpoint. However, your customers are complaining about the unavailability of the API endpoint. Upon checking, you noticed that the Gateway URL is throwing Service Unavailable error. You need to check the backend latency and backend responses when this error started last night. What should you do to get this data? (Choose the best answer.)

- A. Check with the application owner and search the log file for the container to get the metrics from the log file.
- B. Go to Governance Menu and click on Audit to see the Audit log for the API Gatewa
- C. Filter it using Start and End date with a 503 response status.
- D. Go to Developer Services and click on API Gatewa
- E. Go to the detail page of the gateway and select Metric
- F. Change the Start and End time to filter the metrics.
- G. Go to Monitoring and click on Service Metric
- H. Choose the Metric Namespace as oci_apigateway. Change the Start and End time accordingl
- I. Add a Dimension and select httpStatusCode: 503. Check the backend latency and backend responses metric.

Answer: D

Explanation:

<https://medium.com/oracledevs/using-oci-monitoring-healthchecks-to-schedule-execution-of-serverless-function>

NEW QUESTION 19

Multiple departments In your company use a shared Oracle Cloud Infrastructure (OCI) tenancy to Implement their projects. You are in charge of managing the cost of OCI resources in the tenancy and need to obtain better Insights Into department's usage. Which three options can you implement together to accomplish this?

- A. Create a budget that matches your commitment amount and an alert at 100 percent of the forecast
- B. Set up a consolidated budget tracking lags to analyze costs in ,1 granular manner
- C. Set up different compartments for each department then track and analyze cost per compartment
- D. Use the billing cost tracking report to analyze costs
- E. Set up a tag default that automatically applies tags to all specified resources created In a compartment then use these tags for cost analysis.

Answer: ACE

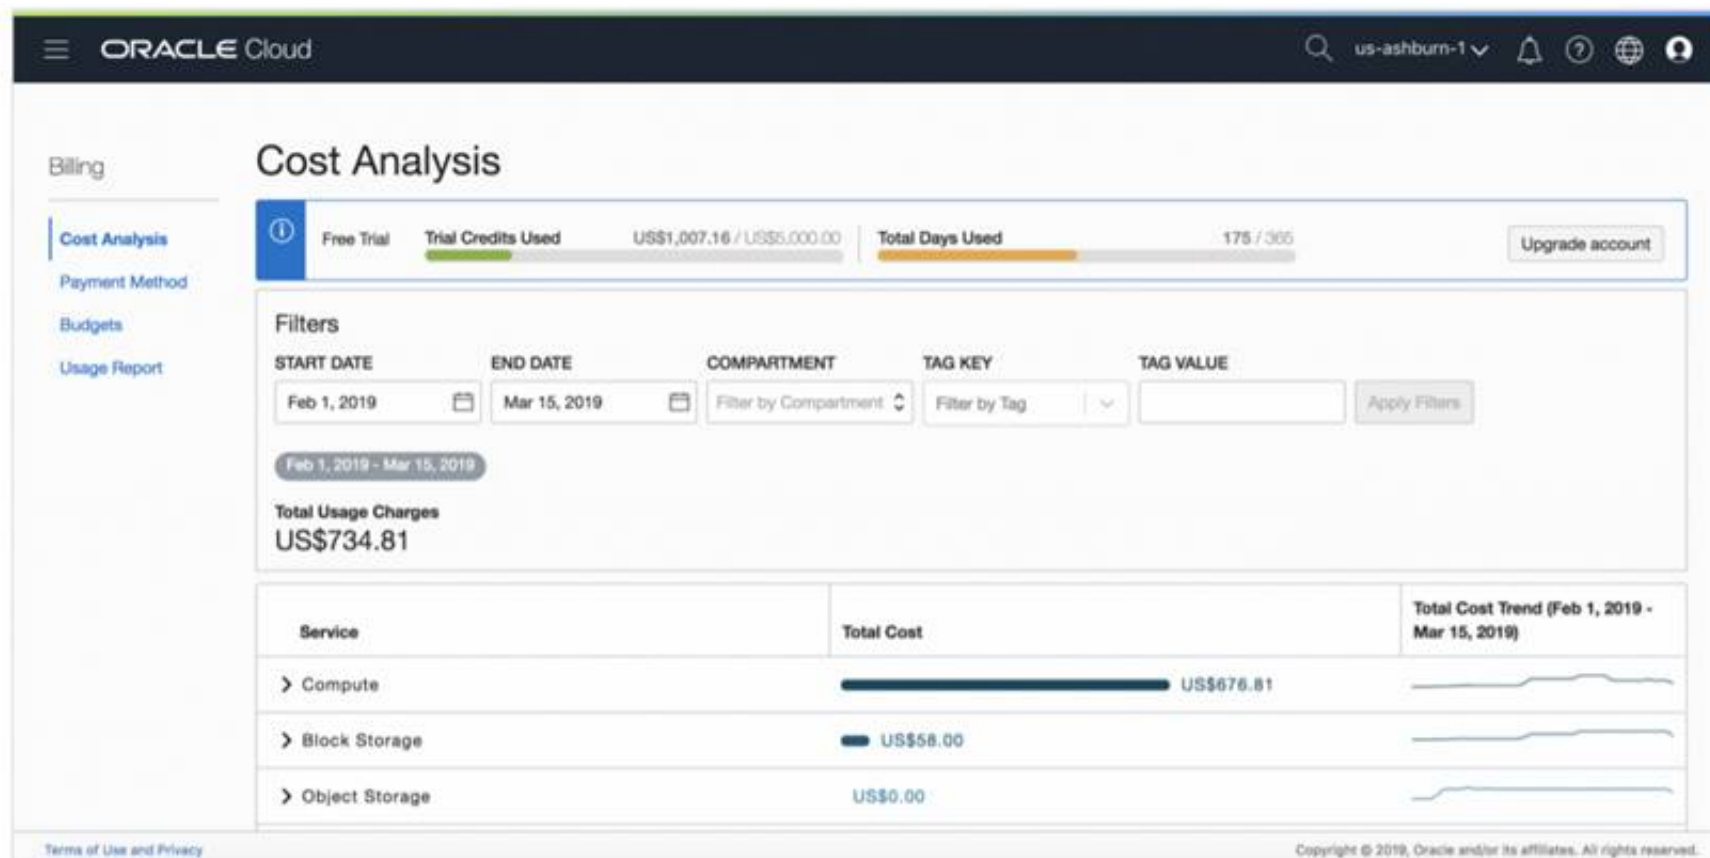
Explanation:

budgets

You can use budgets to track costs in your tenancy. After creating a budget for a compartment, you can set up alerts that will notify you if a budget is forecast to be exceeded or if spending surpasses a certain amount.

OCI Cost Analysis

- Visualization tools Help understand spending patterns at a glance
- Filter costs by Date, Tags and Compartments
- Trend lines show how spending patterns are changing
- To use Cost Analysis you must be a member of the Administrators group



<https://www.oracle.com/a/ocom/docs/cloud/ops-billing-100.pdf>

NEW QUESTION 22

You are working as a solution architect with a global automotive provider who is looking to create a multi-cloud solution. They want to run their application tier in Microsoft Azure while utilizing the Oracle DB Systems in the Oracle Cloud Infrastructure (OCI). What is the most-fault tolerant and secure solution for this customer? (Choose the best answer.)

- A. Deploy the Oracle database system into a public subnet in your VCN and assign a public IP address. Connect your application tier running in Azure to the public IP address of the database system over the internet.
- B. Create a FastConnect virtual circuit with Microsoft Azure as the provider to establish a private interconnect between the application tier running in the Azure Virtual Network and the OCI VCN that contains the Oracle Databases.
- C. Create an encrypted, Virtual Private Network connection between the Microsoft Azure Virtual Network that contains the application tier and the OCI Virtual Cloud Network (VCN) that contains the Oracle Databases.
- D. Use an OCI Virtual Cloud Network remote peering connection to create a remote network connection between the application tier running in Microsoft Azure Virtual Network and Oracle Databases running in the OCI Virtual Cloud Network (VCN).

Answer: B

Explanation:

<https://docs.oracle.com/en/solutions/learn-azure-oci-interconnect/index.html#GUID-FBE38C70-A4CF-40C5-A3>

NEW QUESTION 23

A fast growing E-commerce company has deployed their online shopping application on Oracle Cloud Infrastructure. The application was deployed on compute instances with Autoscaling configuration for application servers fronted by a load balancer and OCI Autonomous Transaction Processing (ATP) in the backend. In order to promote their e-commerce platform 50% discount was announced on all the products for a limited period. During the day 1 of promotional period it was observed that the application is running slow and company's hotline is flooded with complaints. What could be two possible reasons for this situation?

- A. The health check on some of the backend servers has failed and the load balancer has taken those servers temporarily out of rotation.
- B. As part of Autoscaling, the load balancer shape has dynamically changed to a larger shape to handle more incoming traffic and the system was slow for a short time during this change.
- C. The health check on some of the backend servers has failed and the load balancer was rebooting these servers.
- D. Autoscaling has already scaled to the maximum number of instances specified in the configuration and there is no room for scaling further.

Answer: AD

NEW QUESTION 28

You have multiple IAM users who launch different types of compute Instances and block volumes every day. As a result, your Oracle cloud Infrastructure (OCI) tenancy quickly hit the service limit and you can no longer create any new instances. As you are cleaning up environment, you notice that the majority of the Instances and block volumes are untagged. Therefore, It is difficult to pinpoint the owner of these resources verify if they are safe to terminate. Because of this, your company has issued a new mandate, which requires adding compute instances. Which option is the simplest way to implement this new requirement?

- A. Create a policy to automatically tag a resource with the user name.
- B. Create a policy using IAM requiring users to tag specific resource
- C. This will allow a user to launch compute instances only if certain tags were defined.
- D. Create tag variables to automatically tag a resource with the user name.
- E. Create a default tag for each compartment, which ensure that appropriate tags are applied at resource creation
- F. Create tag variables for each compartment to automatically tag a resource with the user name.

Answer: C

Explanation:

Tag Variables

You can use a variable to set the value of a defined tag. When you add the tag to a resource, the variable resolves to the data it represents. You can use tag variables in defined tags and default tags.

Supported Tag Variables

The following tag variables are supported.

`${iam.principal.name}` The name of the principal that tagged the resource

`${iam.principal.type}` The type of principal that tagged the resource.

`${oci.datetime}` The date and time that the tag was created. Consider the following example:

Operations.CostCenter=" `${iam.principal.name}` at `${oci.datetime}` "

Operations is the namespace, CostCenter is the tag key, and the tag value contains two tag

variables `${iam.principal.name}` and `${oci.datetime}` . When you add this tag to a resource, the variable resolves to your user name (the name of the principal that applied the tag) and a time date stamp for when you added the tag.

user_name at 2019-06-18T18:00:57.604Z

The variable is replaced with data at the time you apply the tag. If you later edit the tag, the variable is gone and only the data remains. You can edit the tag value in all the ways you would edit any other tag value. To create a tag variable, you must use a specific format.

`${<variable>}` Type a dollar sign followed by open and close curly brackets. The tag variable goes between the curly brackets. You can use tag variables with other tag variables and with string values. Tag defaults let you specify tags to be applied automatically to all resources, at the time of creation, in a specific compartment. This feature allows you to ensure that appropriate tags are applied at resource creation without requiring the user who is creating the resource to have access to the tag namespaces.

<https://docs.cloud.oracle.com/en-us/iaas/Content/Tagging/Tasks/managingtagdefaults.htm>

NEW QUESTION 33

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?

- A. 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.
- B. 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.
- C. 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.
- D. Stream all the logs and cloud events of Events service to Oracle Streaming Servic
- E. Build a client process that will apply heuristics on the logs and store them in an Object Storage.

Answer: D

Explanation:

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.

NEW QUESTION 38

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.

Answer: C

Explanation:

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

NEW QUESTION 43

You want to automate the processing of new image files to generate thumbnails. The expected rate is 10 new files every hour.

Which of the following is the most cost effective option to meet this requirement in Oracle Cloud Infrastructure (OCI)?

- A. Upload all files to an Oracle Streaming Service (OSS) strea
- B. Setup a cron job to invoke a function in Oracle Functions to fetch data from the strea
- C. Invoke another function to process the image files and generate thumbnails.Store thumbnails in another OSS stream.
- D. Upload files to an OCI Object storage bucke

- E. Every time a file is uploaded, an event is emitted
- F. Write a rule to filter these events with an action to trigger a function in Oracle Function
- G. The function processes the image in the file and stores the thumbnails back in an Object storage bucket.
- H. Build a web application to ingest the files and save them to a NoSQL Database
- I. Configure OCI Events service to trigger a notification using Oracle Notification Service (ONS). ONS invokes a custom application to process the image files to generate thumbnail
- J. Store thumbnails in a NoSQL Database table.
- K. Upload files to an OCI Object storage bucket
- L. Every time a file is uploaded, trigger an event with an action to provision a compute instance with a cloud-init script to access the file, process it and store it back in an Object storage bucket
- M. Terminate the instance using Autoscaling policy after the processing is finished.

Answer: B

NEW QUESTION 46

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.

Answer: D

NEW QUESTION 48

Given this compartment structure:



You are managing a compute instance that currently resides in the Compute compartment. The Virtual Cloud Network (VCN) into which the compute instance was originally deployed, also resides in this compartment. To support a project-related task, you need to move just the compute instance to the SysTest-Team compartment. You log into your Oracle Cloud Infrastructure (OCI) account and use the Move Resource option to place the compute instance in the new compartment.

What will be the result of your attempt to move the compute instance to the new compartment? (Choose the best answer.)

- A. The move will be successful
- B. The compute instance's public and private IP addresses will stay the same. The compute instance will remain associated with the VCN from the source compartment.
- C. The move will fail and you will be prompted to move the VCN first
- D. Once VCN is moved to the target compartment, the compute instance can be moved.
- E. After moving the compute instance, you must move the compute instance VNIC as a separate action. The public and private IP addresses of the instance will remain unchanged and it will still be associated with the VCN from the source compartment.
- F. The move will be successful
- G. However, the compute instance's public and private IP addresses will change, and it will be associated to the first VCN that was created in the new, target compartment.

Answer: C

NEW QUESTION 50

A civil engineering company is running an online portal in which engineers can upload their constructions photos, videos, and other digital files.

There is a new requirement for you to implement: the online portal must offload the digital content to an Object Storage bucket for a period of 72 hours. After the provided time limit has elapsed, the portal will hold all the digital content locally and wait for the next offload period.

Which option fulfills this requirement?

- A. Create a pre-authenticated URL for the entire Object Storage bucket to read and list the content with an expiration of 72 hours.
- B. Create a pre-authenticated URL for each object that is uploaded to the Object Storage bucket with an expiration of 72 hours.
- C. Create a Dynamic Group with matching rule for the portal compute instance and grant access to the Object Storage bucket for 72 hours.
- D. Create a pre-authenticated URL for the entire Object Storage bucket to write content with an expiration of 72 hours.

Answer: D

Explanation:

Pre-authenticated requests provide a way to let users access a bucket or an object without having their own credentials, as long as the request creator has permission to access those objects.

For example, you can create a request that lets operations support user upload backups to

a bucket without owning API keys. Or, you can create a request that lets a business partner update shared data in a bucket without owning API keys.

When creating a pre-authenticated request, you have the following options:

You can specify the name of a bucket that a pre-authenticated request user has write access to and can upload one or more objects to.

You can specify the name of an object that a pre-authenticated request user can read from, write to, or read from and write to.

Scope and Constraints

Understand the following scope and constraints regarding pre-authenticated requests:

Users can't list bucket contents.

You can create an unlimited number of pre-authenticated requests. There is no time limit to the expiration date that you can set.

You can't edit a pre-authenticated request. If you want to change user access options in response to changing requirements, you must create a new pre-authenticated request.

The target and actions for a pre-authenticated request are based on the creator's permissions. The request is not, however, bound to the creator's account login credentials. If the creator's login credentials change, a pre-authenticated request is not affected.

You cannot delete a bucket that has a pre-authenticated request associated with that bucket or with an object in that bucket.

NEW QUESTION 51

A customer has a Virtual Machine instance running in their Oracle Cloud Infrastructure tenancy. They realized that they wrongly picked a smaller shape for their compute instance. They are reaching out to you to help them fix the issue.

Which of the below options is best recommended to suggest to the customer?

- A. Delete the running instance and spin up a new instance with the desired shape.
- B. Change the shape of instance without reboot, but stop all the applications running on instance beforehand to prevent data corruption.
- C. Change the shape of the virtual machine instance using the Change Shape feature available in the console.
- D. OCI doesn't allow such an operation.

Answer: C

Explanation:

You can change the shape of a virtual machine (VM) instance without having to rebuild your instances or redeploy your applications. This lets you scale up your Compute resources for increased performance, or scale down to reduce cost.

When you change the shape of an instance, you select a different processor, number of cores, amount of memory, network bandwidth, and maximum number of VNICs for the instance. The instance's public and private IP addresses, volume attachments, and VNIC attachments remain the same.

NEW QUESTION 52

You are helping a customer troubleshoot a problem. The customer has several Oracle Linux servers in a private subnet within a Virtual Cloud Network (VCN). The servers are configured to periodically communicate to the Internet to get security patches for applications installed on them.

The servers are unable to reach the Internet. An Internet Gateway has been deployed in the public subnet in the VCN and the appropriate routes are configured in the Route Table associated with the public subnet.

Based on cost considerations, which option will fix this issue?

- A. Create a Public Load Balancer in front of the servers and add the servers to the Backend Set of the Public Load Balancer.
- B. Create another Internet Gateway and configure it as route target for the private subnet.
- C. Implement a NAT instance in the public subnet of the VCN and configure the NAT instance as the route target for the private subnet.
- D. Create a NAT gateway in the VCN and configure the NAT gateway as the route target for the private subnet.

Answer: D

NEW QUESTION 53

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 applicatio
- B. Deploy a 2 node VM DB Systems with Oracle RAC enable
- C. Import the on-premises database to OCI VM DB Systems using Oracle Data Pump and then enable automatic backups.
- D. Launch a compute instance for both the NGINX application server and the database serve
- E. Attach block volumes on the database server compute instance and enable backup policy to backup the block volumes.
- F. Launch a compute instance and run an NGINX server to host the applicatio
- G. Deploy Exadata Quarter Rack, enable automatic backups and import the database using Oracle Data Pump.
- H. Launch a compute instance and run an NGINX server to host the applicatio
- I. Deploy a 2 node VM DB Systems with Oracle RAC enable
- J. Setup Oracle GoldenGate to synchronize data from their on-premises database to OCIVM Databas
- K. 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 databas
- L. Enable automatic backups for the OCIVM database and then cutover the application from on-premises to OCI.

Answer: D

NEW QUESTION 54

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" } }, { "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

Answer: A

Explanation:

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.

NEW QUESTION 57

You have deployed a web application targeting a global audience across multiple Oracle Cloud Infrastructure (OCI) regions.

You decide to use Traffic Management Geo-Location based Steering Policy to serve web requests to users from the region closest to the user. Within each region you have deployed a public load balancer with 4 servers in a backend set. During a DR test disable all web servers in one of the regions however, traffic Management does not automatically direct all users to the other region.

Which two are possible causes?

- A. You did not setup a Route Table associated with load Balancer's subnet
- B. You did not setup an HTTP Health Check associated with Load Balancer public IP in the disabled region.
- C. Rather than using Geo-Location based Steering Policy, you should use Failover Policy Type to serve traffic.
- D. One of the two working web servers in the other region did not pass its HTTP health check
- E. You did not correctly setup the Load Balancer HTTP health check policy associated with backend set

Answer: BE

Explanation:

Managing Traffic Management GEOLOCATION Steering Policies

Geolocation steering policies distribute DNS traffic to different endpoints based on the location of the end user. Customers can define geographic regions composed of originating continent, countries or states/provinces (North America) and define a separate endpoint or set of endpoints for each region.

The Health Checks service allows you to monitor the health of IP addresses and hostnames, as measured from geographic vantage points of your choosing, using HTTP and ping probes. After configuring a health check, you can view the monitor's results. The results include the location from which the host was monitored, the availability of the endpoint, and the date and time the test was performed.

Also you can Combine Managing Traffic Management GEOLOCATION Steering Policies with Oracle Health Checks to fail over from one region to another

The Load Balancing service provides health status indicators that use your health check policies to report on the general health of your load balancers and their components.

if you misconfigure the health check Protocol between the Load balancer and backend set that can lead to not get an accurate response as example below

If you run a TCP-level health check against an HTTP service, you might not get an accurate response. The TCP handshake can succeed and indicate that the service is up even when the HTTP service is not configured or having other issues. Although the health check appears good customers might experience transaction failures.

NEW QUESTION 62

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 Groups' 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. The Identity and Access Management policy allowing backup failed to move when the compartment was moved.
- B. You are exceeding your volume group backup quota configured.
- C. 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.
- D. A compute instance with multiple block volumes attached cannot move when a compartment is moved.

Answer: A

NEW QUESTION 64

You are developing a Serverless function for your company's IoT project. This function should access Oracle Cloud Infrastructure (OCI) Object Storage to store some files. You choose Oracle Functions to deploy this function on OCI. However, your security team doesn't allow you to carry any API Token or RSA Key to authenticate the function against the OCI API to access the Object Storage.

What should you do to get this function to access OCI Object Storage without carrying any static authentication files? (Choose the best answer.)

- A. Set up a Dynamic Group using the format below: `ALL {resource.type = 'fnfunc', resource.compartment.id = 'ocidl.compartment.oc1..aaaaaaa23_____smwa' }` Create a policy using the format below to give access to OCI Object Storage:
- ```
allow dynamic-group acme-func-dyn-grp to manage objects in compartment acme-storage-compartment where all {target.bucket.name= 'acme-functions-bucket' }
```
- Include a call to a "resource principal provider" in your function code as below: `signer = oci.auth.signers.get_resource_principals_signer()`
- B. Add these two policy statements for your compartment and then include a call to a "resource principal provider" in your function code:
- ```
Allow group acme-functions-developers to inspect repos in tenancy
Allow group acme-functions-developers to manage repos in tenancy where all
{target.repo.name=/acme-web-app*/}
```
- C. There is no way that you can access the OCI resources from a running function.
- D. Add these two policy statements for your compartment to give your function automatic access to all other OCI resources:
- ```
Allow group <group-name> to manage fn-app in compartment <compartment-name>
Allow group <group-name> to manage fn-function in compartment <compartment-name>
```

- A. Option A  
B. Option B  
C. Option C  
D. Option D

**Answer:** A

**Explanation:**

<https://blogs.oracle.com/cloud-infrastructure/getting-started-with-oracle-functions-and-object-storage>

**NEW QUESTION 68**

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 OC  
B. Then to use OCI CLI command to upload database backups to OCI Object Storage Archive tier as the final destination.  
C. Use OCI Storage Gateway to transfer the backup files to OCI Object Storage Archive tier as the final destination.  
D. Setup a FastConnect connection between on-premises data center and OC  
E. Then to use OCI CLI command to upload database backups to OCI Object Storage Standard tier as the final destination.  
F. Use OCI Storage Gateway to transfer the backup files to OCI Object Storage Standard tier as the finaldestination.

**Answer:** D

**NEW QUESTION 71**

Your organization is planning on using Oracle Cloud Infrastructure (OCI) File Storage Service (FSS). You will be deploying multiple compute instance in Oracle Cloud Infrastructure (OCI) and mounting the file system to these compute instances. The file system will hold payment data processed by a Database instance and utilized by compute instances to create a overall inventory report. You need to restrict access to this data for specific compute instances and must be allowed/blocked per compute instance's CIDR block.

Which option can you use to secure access?

- A. Use stateless Security List rule to restrict access from known IP addresses only.  
B. Create a new VCN security list, choose SOURCE TYPE as Service and SOURCE SERVICE as FS  
C. Add stateless ingress and egress rules for specific P address and CIDR blocks.  
D. Use 'Export option' feature of FSS to restrict access to the mounted file systems.  
E. Create and configure OCI Web Application Firewall service with built in DNS based intelligent routing.

**Answer:** C

**Explanation:**

NFS export options enable you to create more granular access control than is possible using just security list rules to limit VCN access. You can use NFS export options to specify access levels for IP addresses or CIDR blocks connecting to file systems through exports in a mount target. Access can be restricted so that each client's file system is inaccessible and invisible to the other, providing better security controls in multi-tenant environments.

Using NFS export option access controls, you can limit clients' ability to connect to the file system and view or write data. For example, if you want to allow clients to consume but not update resources in your file system, you can set access to Read Only. You can also reduce client root access to your file systems and map specified User IDs (UIDs) and Group IDs (GIDs) to an anonymous UID/GID of your choice. For more information about how NFS export options work with other security layers

**NEW QUESTION 74**

An online registration system is currently hosted on one large Oracle Cloud Infrastructure (OCI) Bare metal compute Instance with attached block volume to store of the users' data. The registration system accepts the information from the user, including documents and photos then performs automated verification and processing to check if the user is eligible for registration.

The registration system becomes unavailable at times when there is a surge of users using the system the existing architecture needs improvement as it takes a long time for the system to complete the processing and the attached block volumes are not large enough to use data being uploaded by the users.

Which is the most effective option to achieve a highly scalable solution?

- A. Attach more Block volumes as the data volume increases, use Oracle Notification Service (ONS) to distribute tasks to a pool of compute instances working in parallel, and Auto Scaling to dynamically size the pool of instances depending on the number of notifications received from the Notification Service. Use Resource Manager stacks to replicate your architecture to another region.
- B. Change your architecture to use an OCI Object Storage standard tier bucket, replace the single bare metal instance with a Oracle Streaming Service (OSS) to ingest the incoming requests and distribute the tasks to a group of compute instances with Auto Scaling
- C. Upgrade your architecture to use a pool of Bare metal servers and configure them to use their local SSDs for faster data access. Set up Oracle Streaming Service (OSS) to distribute the tasks to the pool of Bare metal instances with Auto Scaling to dynamically increase or decrease the pool of compute instances depending on the length of the Streaming queue.
- D. Upgrade your architecture to use more Block volumes as the data volume increases
- E. Replace the single bare metal instance with a group of compute instances with Auto Scaling to dynamically increase or decrease the compute instance pools depending on the traffic.

**Answer:** D

#### NEW QUESTION 79

Which of the below options for private access to services within Oracle Cloud Infrastructure (OCI) is NOT valid?

- A. You cannot use the private endpoint for hosts in the on-premises network.
- B. Traffic from an OCI compute instance going through a Service Gateway to Object Storage is routed without being sent over the internet.
- C. You can enable private access to certain services within OCI from your Virtual Cloud Network by using either a private endpoint or a service gateway.
- D. The private endpoint gives hosts within your Virtual Cloud Network access to a given service within Oracle Cloud Infrastructure.

**Answer:** A

#### NEW QUESTION 82

Many development engineers are deploying new instances as part of their projects in Oracle Cloud Infrastructure tenancy, but majority of these instances have not been tagged. You as an administrator of this tenancy want to enforce tagging to identify owners who are launching these instances.

Which option below should be used to implement this requirement?

- A. Create a predefined tag with tag variables to automatically tag a resource with username.
- B. Create a default tag for each compartment which ensures appropriate tags are allowed at resource creation.
- C. Create tag variables for each compartment to automatically tag a resource with user name.
- D. Create an IAM policy to automatically tag a resource with the username.

**Answer:** A

#### NEW QUESTION 86

Which of the below options is true regarding Oracle Cloud Infrastructure's load balancing service?

- A. You can dynamically change the load balancer shape to handle more incoming traffic.
- B. The public load balancer applies a floating public IP address to the primary load balancer.
- C. When you create a private load balancer, the service requires 2 or more subnets to host both the primary and standby load balancers.
- D. A public load balancer is Availability Domain specific in scope.

**Answer:** B

#### NEW QUESTION 90

A global retailer has decided to re-design its e-commerce platform to have a micro-services architecture. They would like to decouple application architecture into smaller, independent services using Oracle Cloud Infrastructure (OCI). They have decided to use both containers and serverless technologies to run these application instances.

Which option should you recommend to build this new platform?

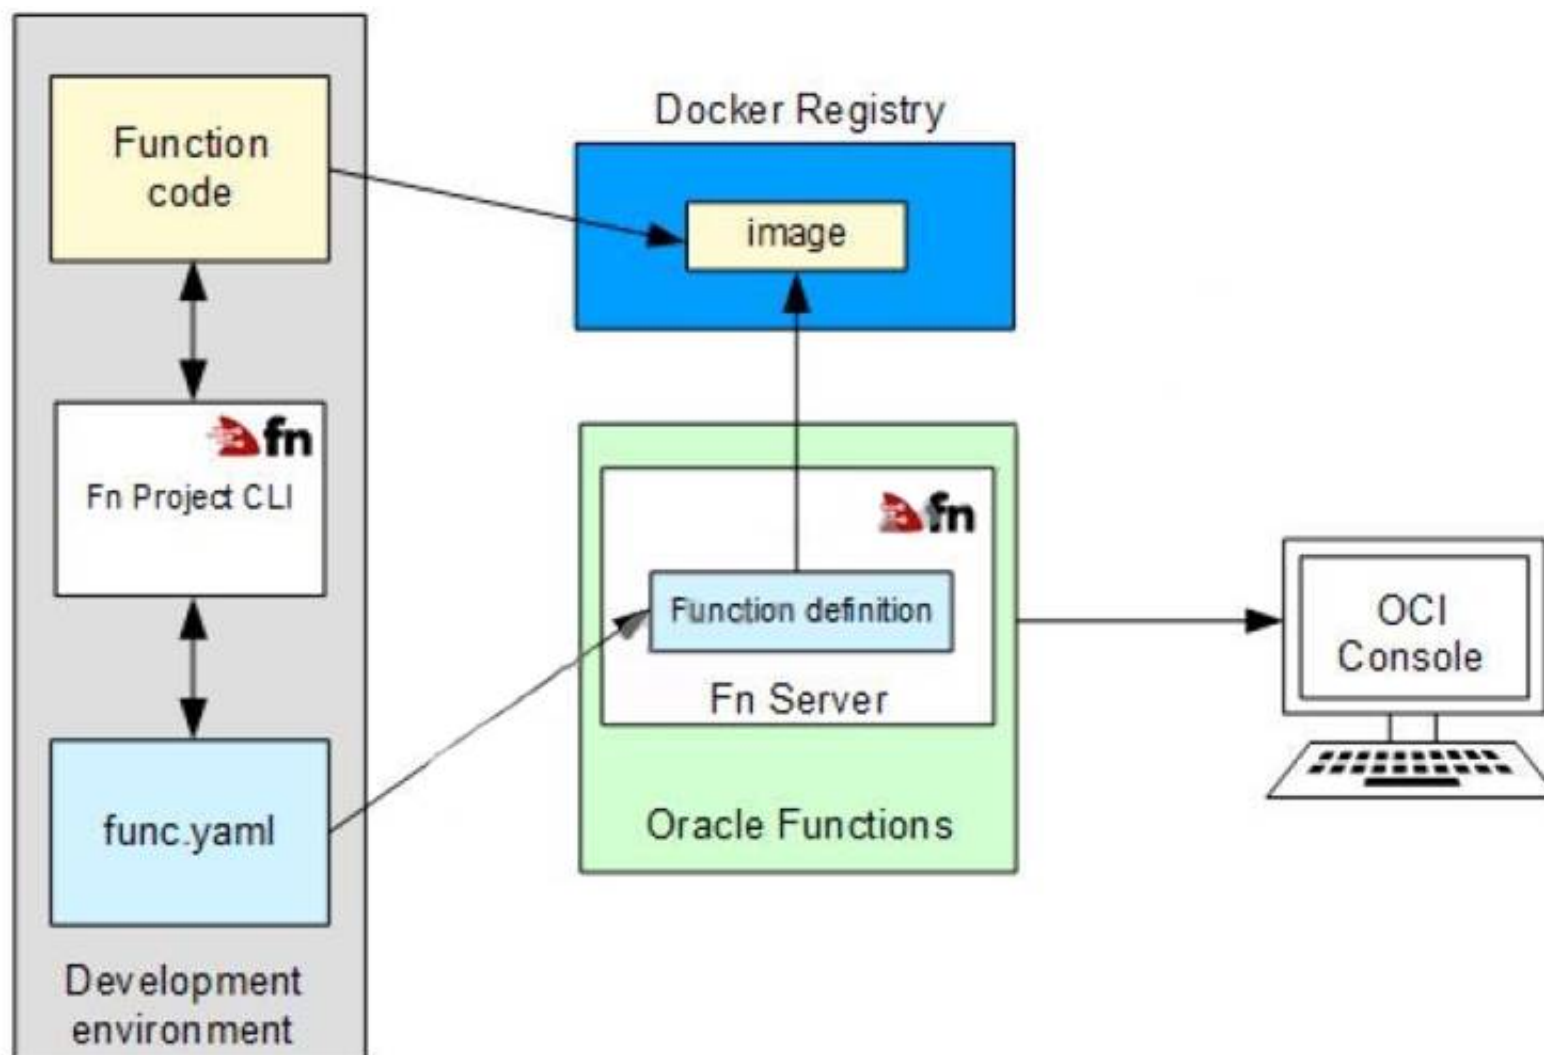
- A. Install a Kubernetes cluster on OCI and use OCI event service.
- B. Use Oracle Container Engine for Kubernetes, OCI Registry and OCI Functions.
- C. Use OCI Resource Manager to automate compute instances provisioning and use OCI Streaming service.
- D. Use OCI functions, OCI object storage and OCI event service.

**Answer:** B

#### Explanation:

Oracle Functions is a fully managed, multi-tenant, highly scalable, on-demand, Functions-as-a-Service platform. It is built on enterprise-grade Oracle Cloud Infrastructure and powered by the Fn Project open source engine. Use Oracle Functions (sometimes abbreviated to just Functions) when you want to focus on writing code to meet business needs.





Oracle Cloud Infrastructure Container Engine for Kubernetes is a fully-managed, scalable, and highly available service that you can use to deploy your containerized applications to the cloud. Use Container Engine for Kubernetes (sometimes abbreviated to just OKE) when your development team wants to reliably build, deploy, and manage cloud-native applications. You specify the compute resources that your applications require, and Container Engine for Kubernetes provisions them on Oracle Cloud Infrastructure in an existing OCI tenancy.

#### NEW QUESTION 92

As a part of migration exercise for an existing on premises application to Oracle Cloud Infrastructure (OCI), you are 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 split this file into 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 CreateMultiPartUpload REST 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

**Answer: AD**

#### Explanation:

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.

#### NEW QUESTION 93

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 24x7.

Which additional configuration will you need to implement to make sure WAF is protecting my web application 24x7?

- 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 new vulnerabilities and mitigations

**Answer: C**

#### Explanation:

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)

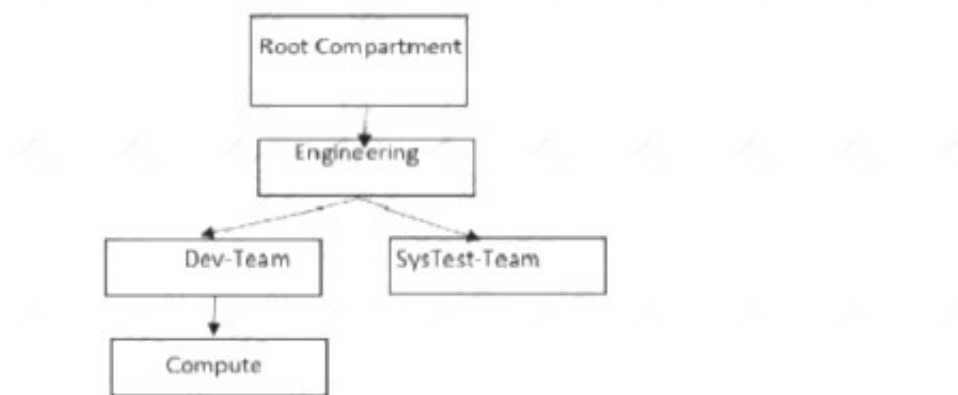
A DDoS attack is an often intentional attack that consumes an entity's resources, usually using a large number of distributed sources. DDoS can be categorized into either Layer 7 or Layer 3/4 (L3/4)

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.

#### NEW QUESTION 97

Give this compartment structure:



You want to move a compute instance that is in 'Compute' compartment to 'SysTes-Team'.

You login to your Oracle Cloud Infrastructure (OCI)account and use the 'Move Resource' option. What will happen when you attempt moving the compute resource?

- A. The move will be successful though Compute Instance and its Public and Private IP address will stay the sam
- B. The Compute instance VNIC will need to be moved separatel
- C. The Compute instance will still be associated with the original VCN.
- D. The move will fail and you will be prompted to move the VCN firs
- E. Once VCN is moved to the target compartment, the Compute instance can be moved.
- F. The move will be successful though Compute Instance Public and Private IP address changed, and it will be associated to the VCN in target compartment.
- G. The move will be successful though Compute Instance and its Public and Private IP address will stay the sam
- H. The Compute instance VNIC will still be associated with the original VCN.

**Answer: D**

#### Explanation:

Moving Resources to a Different Compartment

Most resources can be moved after they are created. There are a few resources that you can't move from one compartment to another. Some resources have attached resource dependencies and some don't.

Not all attached dependencies behave the same way when the parent resource moves.

For some resources, the attached dependencies move with the parent resource to the new compartment.

The parent resource moves immediately, but in some cases attached dependencies move asynchronously and are not visible in the new compartment until the move is complete.

For other resources, the attached resource dependencies do not move to the new compartment. You can move these attached resources independently.

You can move Compute resources such as instances, instance pools, and custom images from one compartment to another. When you move a Compute resource to a new compartment, associated resources such as boot volumes and VNICs are not moved.

You can move a VCN from one compartment to another. When you move a VCN, its associated VNICs, private IPs, and ephemeral IPs move with it to the new compartment.

#### NEW QUESTION 98

You are tasked with backing up your data using Oracle Cloud Infrastructure Block Volume service. When you are finalizing your block volume backup schedule, which of the following two are valid

considerations for your backup plan? (Choose Two)

- A. Number of stored backups: How many backups you need to keep available and the deletion schedule for those you no longer need.
- B. Governance: Tagging of backups so you can capture backup related API calls through the Audit service.
- C. Frequency: How often you want to back up your data.
- D. Location: Determine the Object Store Bucket where the backups will be stored.
- E. Encryption: Whether to use your own key to encrypt your volume backups.

**Answer: AC**

#### NEW QUESTION 100

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 (OCI) 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?

- A. 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.
- B. Setup FastConnect virtual circuit as your primary connection, and a second FastConnect virtual circuit as a backup connectio
- C. 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.
- D. Setup FastConnect virtual circuit as your primary connection, and an IPsec VPN as a backup connectio
- E. Use separate edge devices in your on-premises data center for each connectio
- F. From your edge devices, advertise more specific routes through FastConnect virtual circuit, and more specific routes through the backup IPsec VPN path.

- G. Setup IPSec VPN as your primary connection, and a second IPSec VPN as a backup connectio
- H. Use separate edge devices in your on p data center for each connectio
- I. From your edge devices, advertise more specific routes via primary IPSec VP
- J. and less specific rod the backup IPSec VPN.

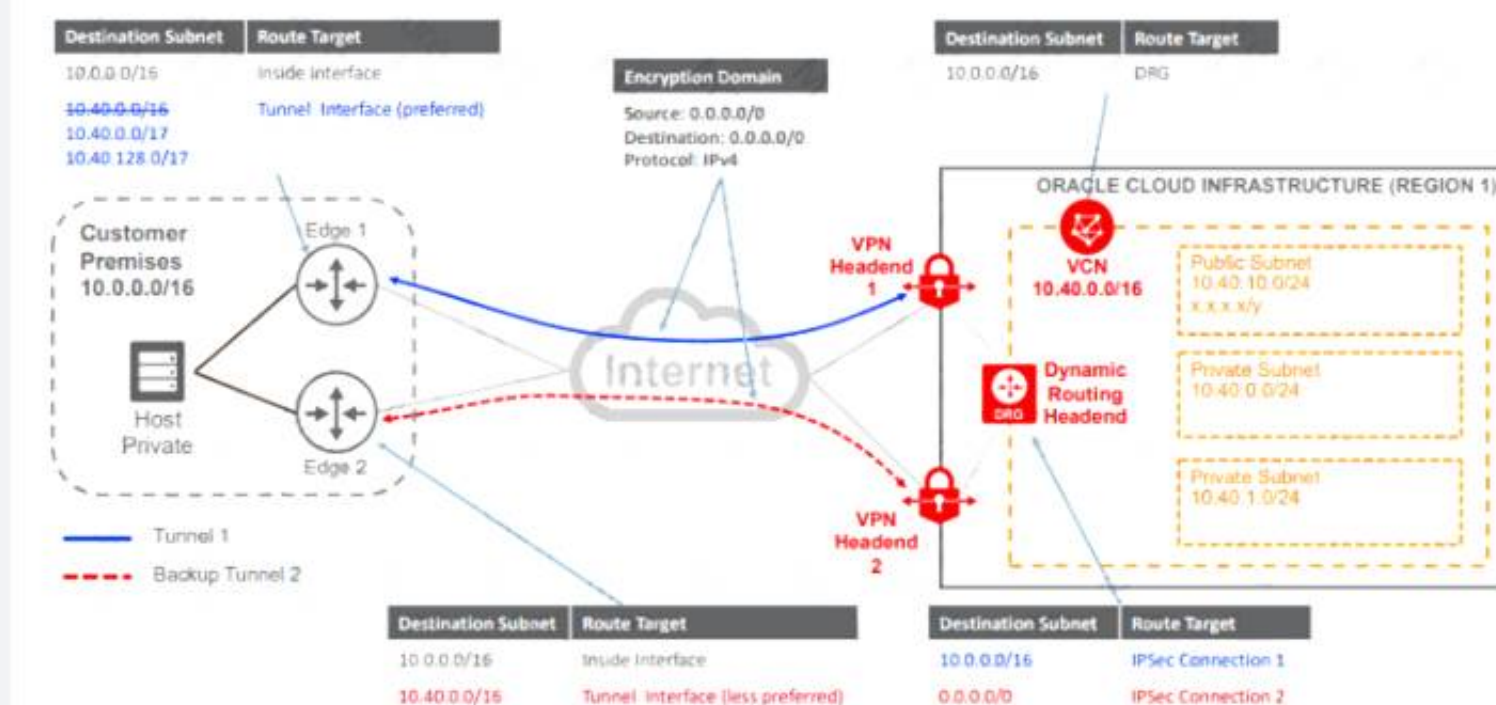
**Answer: D**

**Explanation:**

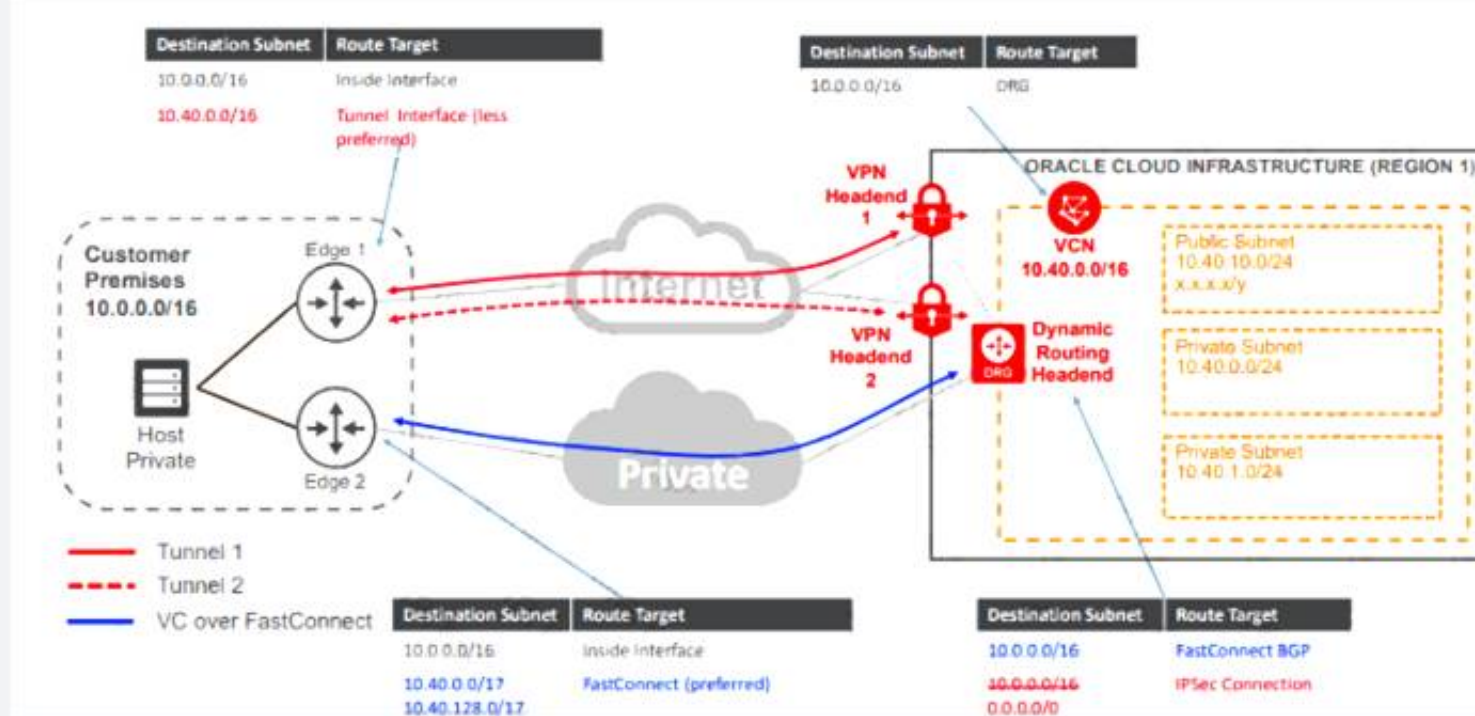
there are two main requirements for this Customer

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

### 1- VPN Connect with a Redundant Customer Edge Device

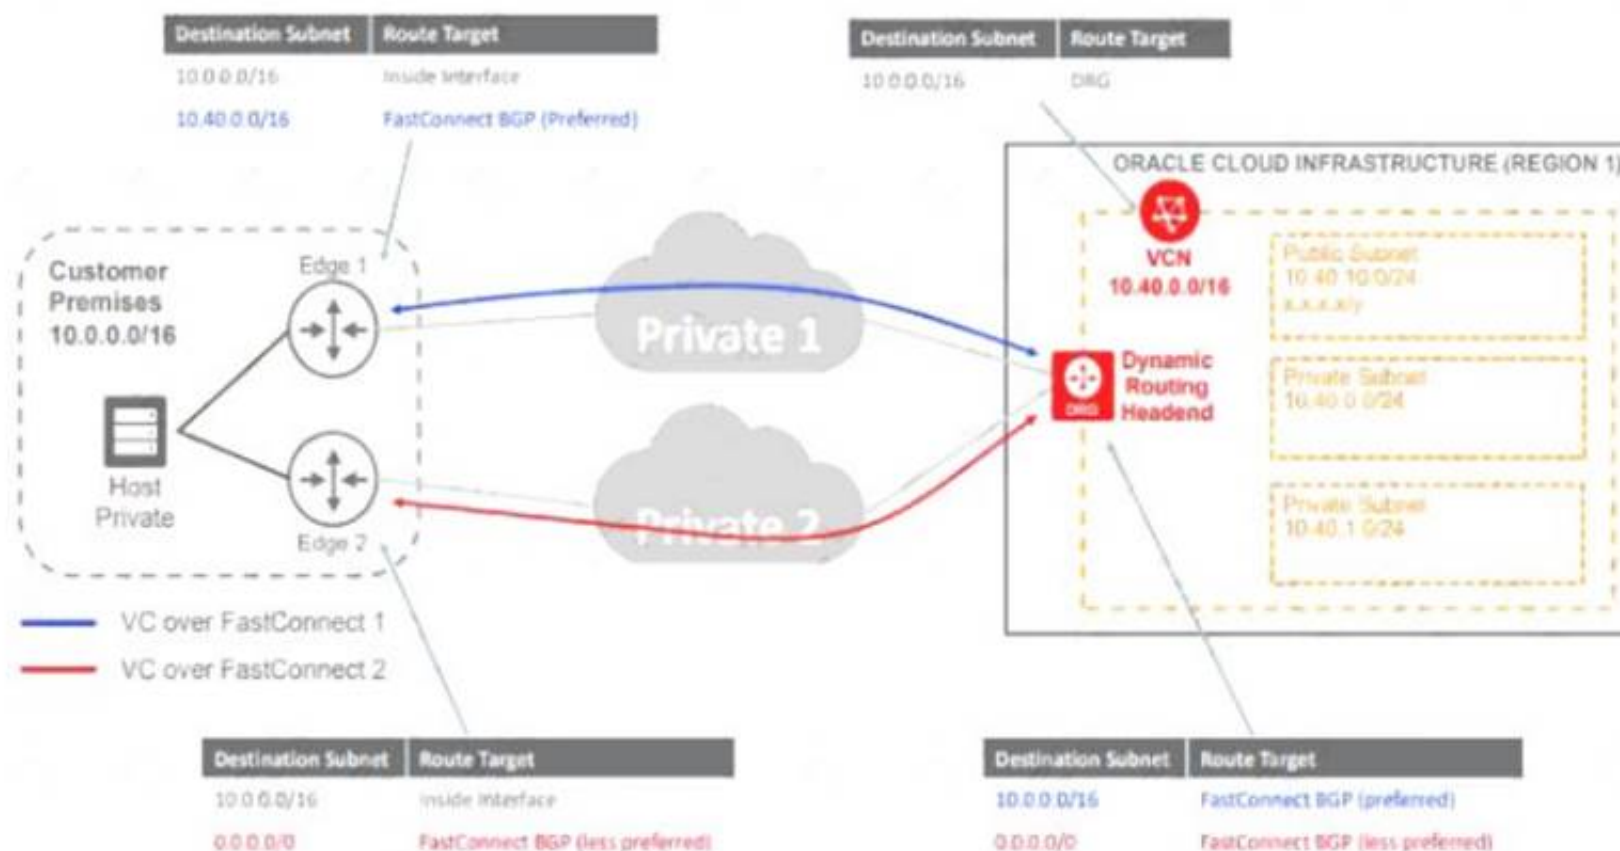


### 2- FastConnect Plus a Single VPN Connect Connection



### 3- Redundant FastConnect





### NEW QUESTION 103

You are a solution architect working with a startup that has decided to move their workload to Oracle Cloud Infrastructure. Since their workload is small, upon architecting, you decide its sufficient to use 8 compute instances to run their workload. The company wants to use a common storage for their instances. So, you propose the idea of attaching a block volume to multiple instances to provide a common storage.

Which of the below option is NOT true for such a solution?

- A. If the block volume is already attached to an instance as read/write non-shareable you can't attach it to another instance until you detach it from the first instance.
- B. Block volumes attached as read-only are configured as shareable by default.
- C. You can delete a block volume from one instance without detaching it from all other instances there by keeping other instance's storage intact.
- D. Once you attach a block volume to an instance as read-only, it can only be attached to other instances as read-only.

**Answer: C**

### NEW QUESTION 105

As part of planning the network design on Oracle Cloud Infrastructure, you have been asked to create an Oracle Cloud Infrastructure Virtual Cloud Network (VCN) with 3 subnets, one in each Availability Domain. Each subnet needs to have a minimum of 64 usable IP addresses.

What is the smallest subnet and VCN size you should use to implement this design? The requirements are static, so no growth is expected.

- A. 122 for the VCN; 124 for the subnets
- B. /23 for the VCN; /25 for the subnets
- C. /24 for the VCN; /24 for the subnets
- D. /22 for the VCN; /25 for the subnets

**Answer: B**

### NEW QUESTION 108

A global media organization is working on a project which lets users upload their videos to the site. After upload is complete, the video should be automatically processed by an AI algorithm. The algorithm will try to recognize certain actions in the videos so that it can be used to show related advertisements in future. The development team wants to focus on writing AI code and not worry about underlying infrastructure for high availability, scalability, security and monitoring. Which Oracle Cloud Infrastructure (OCI) services would meet these requirements?

- A. OCI Object Storage, OCI Events service and OCI Functions.
- B. Oracle Container Engine for Kubernetes, OCI Notifications and OCI Object Storage.
- C. OCI Events, Oracle Container Engine for Kubernetes and OCI Digital Assistant.
- D. OCI Resource Manager, OCI Functions and OCI Events service.

**Answer: A**

### NEW QUESTION 110

A large London based eCommerce company is running Oracle DB System Virtual RAC database on Oracle Cloud Infrastructure (OCI) for their eCommerce application activity. They are launching a new product soon, which is expected to sell in large quantities all over the world.

The application architecture should have minimal cost, no data loss, no performance impacts during the database backup windows and should have minimal downtime.

- A. Launch a new VM RAC database in another availability domain, launch a compute instance, deploy Oracle GoldenGate on it and then configure it to replicate the data from the eCommerce Database over to the new RAC database using GoldenGat
- B. Take backups from the new VM RAC database.
- C. Turn off automated backups from the eCommerce database, implement Oracle Data Guard with the Standby database deployed on another availability domain, take backups from the standby database.
- D. Launch a new VM RAC database in another availability domain, launch a compute instance, deploy Oracle GoldenGate on it and then configure bi-directional replication from the eCommerce Database over to the new VM RAC database using GoldenGat

- E. Take backups from the new VM RAC database.  
F. Turn off automatic backups from the eCommerce database, implement Oracle Active Data Guard with the standby database deployed on another availability domain, and take backups from the standby database.

**Answer: C**

**Explanation:**

Active Data Guard or GoldenGate are used for disaster recovery when fast recovery times or additional levels of data protection are required. And offload queries and backup to standby system.

Oracle GoldenGate to support a disaster recovery site is to have a working bi-directional data flow, from the primary system to the live-standby system and vice versa.

DataGuard and Automatic Backup

You can enable the Automatic Backup feature on a database with the standby role in a Data Guard association.

However, automatic backups for that database will not be created until it assumes the primary role.

**NEW QUESTION 114**

You are working as a solutions architect for an online retail store In Frankfurt which uses multiple compute instance VMs spread among three availability domains In the eu-frankfurt-1 region.

You noticed the website Is having very high traffic, so you enabled autoscaling to sun tee me no f your application but, you observed that one of the availability domains is not receiving any traffic.

What could be wrong In this situation?

- A. Autoscaling only works with single availability domains.  
B. You have to manually acid all three availability domains to your load balancer configuration.  
C. Autoscaling can be enabled for multiple availability domains only in uk-london t region.  
D. Autoscaling is using an Instance Pool configured to create instances in two availability Domains.  
E. You forgot to attach a load balancer to your instance pool configuration.

**Answer: D**

**Explanation:**

Autoscaling lets you automatically adjust the number of Compute instances in an instance pool based on performance metrics such as CPU utilization. This helps you provide consistent performance for your end users during periods of high demand, and helps you reduce your costs during periods of low demand.

you can associate a load balancer with an instance pool. If you do this, when you add an instance to the instance pool, the instance is automatically added to the load balancer's backend set . After the instance reaches a healthy state (the instance is listening on the configured port number), incoming traffic is automatically routed to the new instance.

Instance pools let you provision and create multiple Compute instances based off the same configuration, within the same region.

By default, the instances in a pool are distributed across all fault Domains in a best-effort manner based on capacity. If capacity isn't available in one fault domain, the instances are placed in other fault domains to allow the instance pool to launch successfully.

In a high availability scenario, you can require that the instances in a pool are evenly distributed across each of the fault domains that you specify. When sufficient capacity isn't available in one of the fault domains, the instance pool will not launch or scale successfully, and a work request for the instance pool will return an "out of capacity" error. To fix the capacity error, either wait for capacity to become available,

or use the UpdateInstancePool operation to update the placement configuration (the availability domain and fault domain) for the instance pool.

during create the instance pool you can select the location where you want to place the instances" In the Availability Domain list, select the availability domain to launch the instances in.

If you want the instances in the pool to be placed evenly in one or more fault domains, select the Distribute instances evenly across selected fault domains check box. Then, select the fault domains to place the instances in.

**NEW QUESTION 118**

You are part of a project team working in the development environment created in Oracle Cloud Infrastructure (OCI). You realize that the CIDR block specified for one of the subnets in a Virtual Cloud Network (VCN) is not correct and want to delete the subnet. While deleting you get an error indicating that there are still resources that you must delete first. The error includes the OCID of the VNIC that is in the subnet.

Which of the following action you will take to troubleshoot this issue?

- A. Use OCI CLI to call "network vnic" and "compute vnic-attachment" operations to find out the parent resource of the VNIC.  
B. Use OCI CLI to delete the VNIC first and then delete the subnet.  
C. Use OCI CLI to delete the subnet using -force option.  
D. Copy and paste OCID of the VNIC in the search box of the OCI Console to find out the parent resource of the VNIC.

**Answer: A**

**NEW QUESTION 120**

You are designing the network infrastructure for two application servers: appserver-1 and appserver-2 running in two different subnets inside the same Virtual Cloud Network (VCN) Oracle Cloud Infrastructure (OCI). You have a requirement where your end users will access appserver-1 from the internet and appserver-2 from the on-premises network. The on-premises network is connected to your VCN over a FastConnect virtual circuit.

How should you design your routing configuration to meet these requirements?

- A. Configure a single routing table (Route Table-1) that has two set of rule  
B. One that has route to internet via the internet Gateway and another that propagate specific routes for the on-premise network via the Dynamic Routing Gatewa  
C. Associate the routing table with all the VCN subnets.  
D. Configure a single routing table (Routing Table-1) that has two set of rules: one that has route to internet via the Internet Gateway and another that propagates specific routes for the on-premises network via Dynamic Routing Gateway (DRG). Associate the routing table with the VCN.  
E. Configure two routing tables: Route Table-1 that has a route to internet via the Internet gateway.Associate this route table to the subnet containing appserver-1. Route Table-2 that propagate specific routes for the on-premises network via the Dynamic Routing Gateway (DRG) Associate this route table to subnet containing appserver-2.  
F. Configure two routing table (Route table-1 Route Table-2) that have rule to route all traffic via the Dynamic Routing Gateway (DRG) Associate the two routing tables with all the VCN subnets.

**Answer: C**

**Explanation:**

An internet gateway is an optional virtual router you can add to your VCN to enable direct connectivity to the internet. Resources that need to use the gateway for internet access must be in a public subnet and have public IP addresses. Each public subnet that needs to use the internet gateway must have a route table rule that specifies the gateway as the target. For traffic to flow between a subnet and an internet gateway, you must create a route rule accordingly in the subnet's route table (for example, destination CIDR = 0.0.0.0/0 and target = internet gateway).

Dynamic Routing Gateway (DRG) is A virtual edge router attached to your VCN. Necessary for private peering. The DRG is a single point of entry for private traffic coming in to your VCN, After creating the DRG, you must attach it to your VCN and add a route for the DRG in the VCN's route table to enable traffic flow.

**NEW QUESTION 123**

You are working for a Travel company and your travel portal application is a collection of microservices that run on Oracle Cloud Infrastructure Container Engine for Kubernetes. As per the recent security overview, you have noticed that Oracle has published a newer image of the Operating System used by the worker nodes. You want to make sure that your application doesn't face any downtime but at the same time the worker nodes gets upgraded to the latest version of the Operating System.

What should you do to get this upgrade done without application downtime? (Choose the best answer.)

- A. 1. Shutdown the worker nodes 2. Create a new node pool 3. Manually schedule the pods on the newly built node pool
- B. 1. Create a new node pool using the latest available Operating System image
- C. 2. Run `kubectl cordon <node name>` against all the worker nodes in the old pool to stop any new application pods to get scheduled 3. Run `kubectl drain <node name> --delete-local-data --force --ignore-daemonsets` to evict any Pods that are running 4. Delete the old node pool
- D. 1. Create a new node pool using the latest available Operating System image 2. Run `kubectl taint nodes <all node> role.kubernetes.io/master=` 3. Delete the old node pool
- E. 1. Run `kubectl cordon <node name>` against all the worker nodes in the old pool to stop any new application pods to get scheduled 2. Run `kubectl drain <node name> --delete-local-data --force --ignore-daemonsets` to evict any Pods that are running 3. Download the patches for the new Operating System image 4. Patch the worker nodes to the latest Operating System image

**Answer: B**

**Explanation:**

<https://docs.cloud.oracle.com/en-us/iaas/Content/ContEng/Tasks/contengupgradingk8sworkernode.htm>

**NEW QUESTION 126**

You work for a retail company and they developed a Microservices based shopping application that needs to access Oracle Autonomous Database from the application. As an Architect, you have been tasked to treat all of the application components as Kubernetes native objects, such as the microservices, Oracle Autonomous database, Kubernetes services, etc.

What should you do to make sure that you can use Kubernetes constructs to manage the life cycle of the application components, including Oracle Autonomous Database? (Choose the best answer.)

- A. Create an Oracle Cloud Infrastructure (OCI) Service Gateway and connect to the Oracle Autonomous Database using the private IP address from the microservice.
- B. Provision an Oracle Autonomous Database and then use OCI Service Broker to access the database as a native component to your Kubernetes cluster.
- C. Create a service from the Kubernetes cluster and point to the Oracle Autonomous Database using its FQDN.
- D. Install and secure the OCI Service Broker for Kubernetes
- E. Then provision and bind to the required Oracle Cloud Infrastructure services.

**Answer: D**

**Explanation:**

OCI Service Broker for Kubernetes is an implementation of the Open Service Broker API. OCI Service Broker for Kubernetes is specifically for interacting with Oracle Cloud Infrastructure services from Kubernetes clusters. It includes three service broker adapters to bind to the following Oracle Cloud Infrastructure services: Object Storage Autonomous Transaction Processing Autonomous Data Warehouse

**NEW QUESTION 130**

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

**Answer: C**

**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's effectively equivalent to a virtual independent HSM.

**NEW QUESTION 134**

You are working as a solution architect for a customer in Frankfurt, which uses multiple compute instance VMs spread among three Availability Domains in the



Oracle Cloud Infrastructure (OCI) eu-frankfurt-1 region. The compute instances do not have public IP addresses and are running in private subnets inside a Virtual Cloud Network (VCN). You have set up OCI Autoscaling feature for the compute instances, but find out that instances cannot be auto scaled. You have enabled monitoring on the instances.

What could be wrong in this situation?

- A. You need to assign a reserved public IP address to the compute instances.
- B. You need to set up a Service Gateway to send metrics to the OCI Monitoring service.
- C. Autoscaling only works for instances with public IP addresses.
- D. Autoscaling only works with single availability domains.

**Answer:** B

#### NEW QUESTION 139

You are creating an Oracle Cloud Infrastructure Dynamic Group. To determine the members of this group you are defining a set of matching rules. Which of the following are the supported variables to define conditions in the matching rules? (Choose Two)

- A. instance.compartment.id -the OCID of the compartment where the instance resides.
- B. instance.tenancy.id -the OCID of the tenancy where the instance resides.
- C. tag.<tagnamespace>.<tagkey>.value -the tag namespace and tag key.
- D. iam.policy.id - the OCID of the IAM policy to apply to the group.

**Answer:** AC

#### NEW QUESTION 143

A new international hacktivist group, based in London, launched wide scale cyber attacks including SQL Injection and Cross-Site Scripting (XSS) across multiple websites which are hosted in Oracle Cloud Infrastructure (OCI). As an IT consultant, you must configure a Web Application Firewall (WAF) to protect these websites against the attacks.

How should you configure your WAF to protect the website against those attacks? (Choose the best answer.)

- A. Enable an Access Rule that contains XSS Filters Categories and SQL Filters Categories.
- B. Enable a Protection Rule to block the attacks based on HTTP Headers that contain XSS and SQL strings.
- C. Enable a Protection Rule that contains XSS Filters Categories and SQL Filters Categories.
- D. Enable an Access Rule to block the IP Address range from London.
- E. Enable a Protection Rule to block requests that came from London.

**Answer:** C

#### Explanation:

<https://www.ateam-oracle.com/using-oci-waf-web-application-firewall-with-oracle-e-business-suite#:~:text=The>

#### NEW QUESTION 147

Which of the following features is NOT supported by Oracle Cloud Infrastructure Multi-factor authentication (MFA)?

- A. Only the user can enable MFA for their own account.
- B. Members of the Administrators group can disable MFA for other users.
- C. Users can disable MFA for their own accounts.
- D. Members of the Administrators group can enable MFA for other users.

**Answer:** D

#### NEW QUESTION 151

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

- A. Elapsed time since database was last patched
- B. On-premises host operating system and version
- C. Number of active connections
- D. Data types used in the on-premises database
- E. Top 5 longest running queries

**Answer:** BD

#### Explanation:

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

#### NEW QUESTION 153

You are trying to delete a compartment. The delete operation is failing and you need to troubleshoot the problem.

Which step should NOT be considered when troubleshooting this issue?

- A. Verify that there are no policies in the root compartment that reference the compartment you are trying to delete.
- B. Verify that you have removed all resources from the compartment.
- C. Make sure you have at least one more compartment in your tenancy other than the root compartment.
- D. Search for resources in the compartment for each region that your tenancy is subscribed to.

**Answer:** A

#### NEW QUESTION 157

To serve web traffic for a popular product, your cloud engineer has provisioned four BM.Standard2.52 instances, event spread across two availability domains in the us-asburn-1 region: LoadBalancer is used to deliver the traffic across instances.

After several months, the product grows even more popular and you need additional compute capacity. As a result, an engineer provisioned two additional VM.Standard2.8 instances.

You register the two VM. Standard2. 8 Instances with your load Balancer Backend sot and quickly find that the VM Standard2.8 Instances running at 100% of CPU utilization but the BM.Standard2 .52 instances have significant CPU capacity that's unused.

Which option is the most cost effective and uses instances capacity most effectively?

- A. Configure your Load Balance, with weighted round robin policy to distribute traffic to the compute instances, with more weight assigned to bare metal instances.
- B. Configure Autoscaling instance pool with LoadBalancer to add up to 3 more BM.Standard2.52 Instances when trigger
- C. Shut off VM.Standard2.8 instances.
- D. Route traffic to BM.Standard2.52 and VM Standard2.8 instances directly using DNS and Health Check
- E. Shut off the load Balances.
- F. Configure LoadBalancer with two VM Standard2.8 instances and use Autoscaling Instant pool to add up to two additional VM instance
- G. Shut off BM.Standard2.52 instances.

**Answer:** A

#### Explanation:

Customer have 4 BM.Standard2.52 and After several months he need additional compute capacity customer find The VM Standard2.8 Instances running at 100% of CPU utilization but the BM.Standard2 .52 instances have significant CPU capacity that unused.

so the customer need to check the Load balance policy to make sure the 4 BM and VM is utilize correctly

#### NEW QUESTION 162

Your customer recently ordered for a 1-Gbps Fast Connect connection In ap-tokyo-1 region of Oracle Cloud Infrastructure (OCI). They will us this to one Virtual cloud Network (VCN) in their production (OC1) tenancy and VCN In their development OC1 tenancy

As a Solution Architect, how should yon configure and architect the connectivity between on premises and VCNs In OCI?

- A. Create two private virtual circuits on the FastConnect lin
- B. Create two Dynamic Routing Gateways, one for each VCN
- C. Attach the virtual circuits to the dynamic routing gateways.
- D. You cannot achieve connectivity using single FastConnect link as the production and the development VCNs-are in separate tenancie
- E. Request one more FastConnect connection.
- F. Create a single private virtual circuit over FastConnect and attach fastConnect to either of the VCN's Dynamic Routing Gatewa
- G. Use Remote Peering to peer production and development VCNs.
- H. Create a hub-VCN that uses Dynamic Routing Gateway (DRG) to communicate with on-premises network over FastConnec
- I. Connect the hub-VCN to the production VCN spoke and with development VCN spoke, each peered via their respective local Peering Gateway (LPG)

**Answer:** D

#### Explanation:

There's an advanced routing scenario called transit routing that enables communication between an onpremises network and multiple VCNs over a single Oracle Cloud Infrastructure FastConnect or IPsec VPN.

The VCNs must be in the same region and locally peered in a hub-and-spoke layout. As part of the scenario, the VCN that is acting as the hub has a route table associated with each LPG (typically route tables are associated with a VCN's subnets).



#### NEW QUESTION 165

An online Stock trading application is deployed to multiple Availability Domains in the us phoenix-1 region. Considering the high volume of transactions that the trading application handles, the company has hired you to ensure that the data stored by the application available, and disaster resilient. In the event of failure, the Recovery lime Objective (UK)) must be less than 2 hours to meet regulator requirements.

Which Disaster Recovery strategy should be used to achieve the RTO requirement In the event of system failure?

- A. Configure hourly block volumes backups through the Storage Gateway service.
- B. Configure hourly block volumes backups using the Oracle Cloud Infrastructure (OCI) Command Line Interface (CLI)
- C. Store hourly block volumes backup to NVMe device under a compute instance and generate a custom Image every 5 minutes.
- D. Configure your application to use synchronous master slave data replication between Availability Domains.

**Answer:** B

#### Explanation:

You can use the CLI, REST APIs, or the SDKs to automate, script, and manage volume backups and their lifecycle.

#### Planning Your Backup

The primary use of backups is to support business continuity, disaster recovery, and long-term archiving requirements. When determining a backup schedule, your backup plan and goals should consider the following:

Frequency: How often you want to back up your data.

Recovery time: How long you can wait for a backup to be restored and accessible to the applications that use it. The time for a backup to complete varies on several factors, but it will generally take a few minutes or longer, depending on the size of the data being backed up and the amount of data that has changed since your last backup.

Number of stored backups: How many backups you need to keep available and the deletion schedule for those you no longer need. You can only create one backup at a time, so if a backup is underway, it will need to complete before you can create another one. For details about the number of backups you can store

#### NEW QUESTION 170

A global media organization is working on a project which lets users upload their videos on their site. After upload is complete, the video should be automatically processed by an AI algorithm. The algorithm will try to recognize actions in the videos so that it can be used to show related advertisements in future. The development team wants to focus on writing AI code and don't want to worry about underlying infrastructure for high-availability, scalability, security and monitoring. Which OCI services should you recommend for this project?

- A. Use OCI Events service for triggering automatic processing of video, Oracle Container Engine for Kubernetes (OKE) and OCI Digital Assistant
- B. Use Oracle Container Engine for Kubernetes (OKE) for deployment of AI Code, OCI Notifications and Object Storage
- C. Use OCI Resource Manager to manage the underlying infrastructure, OCI Functions and OCI Events service.
- D. Use Object Storage for storing videos, OCI Events service and OCI Functions

**Answer:** D

#### Explanation:

Oracle Functions is a fully managed, multi-tenant, highly scalable, on-demand, Functions-as-a-Service platform. It is built on enterprise-grade Oracle Cloud Infrastructure and powered by the Fn Project open source engine. Use Oracle Functions (sometimes abbreviated to just Functions) when you want to focus on writing code to meet business needs.

The serverless and elastic architecture of Oracle Functions means there's no infrastructure administration or software administration for you to perform. You don't provision or maintain compute instances, and operating system software patches and upgrades are applied automatically. Oracle Functions simply ensures your app is highly-available, scalable, secure, and monitored. With Oracle Functions, you can write code in Java, Python, Node, Go, and Ruby (and for advanced use cases, bring your own Dockerfile, and Graal VM). You can then deploy your code, call it directly or trigger it in response to events, and get billed only for the resources consumed during the execution.

You can create automation based on state changes for your Oracle Cloud Infrastructure resources by using event types, rules, and actions. When the function is executing inside the container, the function can read from and write to other resources and services running in the same subnet (for example, Database as a Service). The function can also read from and write to other shared resources (for example, Object Storage), and other Oracle Cloud Services.

#### NEW QUESTION 174

After performing maintenance on an Oracle Linux compute instance the system is returned to a running state. You attempt to connect using SSH but are unable to do so. You decide to create an instance console connection to troubleshoot the issue.

Which three tasks would enable you to connect to the console connection and begin troubleshooting?

- A. Use SSH to connect to the public: IP address of the compute Instance and provide the console connection OCID as the username.
- B. edit the Linux boot menu to enable access to console.
- C. Use SSH to connect to the service endpoint of the console connection service
- D. Reboot the compute instance using the Oracle Cloud Infrastructure (OCI) Management Console
- E. Upload an API signing key for console connection authentication.
- F. Stop the compute Instance using the Oracle cloud Infrastructure (OCI) Command Line interface (CLI).

**Answer:** BCD

#### Explanation:

The Oracle Cloud Infrastructure Compute service provides console connections that enable you to remotely troubleshoot malfunctioning instances, such as: An imported or customized image that does not complete a successful boot. A previously working instance that stops responding.

the steps to connect to console and troubleshoot the OS Issue

1- Before you can connect to the serial console you need to create the instance console connection. Open the navigation menu. Under Core Infrastructure, go to Compute and click Instances.

Click the instance that you're interested in. Under Resources, click Console Connections.

Click Create Console Connection.

Upload the public key (.pub) portion for the SSH key. You can browse to a public key file on your computer or paste your public key into the text box.

Click Create Console Connection.

When the console connection has been created and is available, the status changes to ACTIVE.

2- Connecting to the Serial Console

you can connect to the serial console by using a Secure Shell (SSH) connection to the service endpoint of the console connection service

Open the navigation menu. Under Core Infrastructure, go to Compute and click Instances.

Click the instance that you're interested in. Under Resources, click Console Connections.

Click the Actions icon (three dots), and then click Copy Serial Console Connection for Linux/Mac.

Paste the connection string copied from the previous step to a terminal window on a Mac OS X or Linux system, and then press Enter to connect to the console.

If you are not using the default SSH key or ssh-agent, you can modify the serial console connection string to include the identity file flag, `-i`, to specify the SSH key to use. You must specify this for both the SSH

connection and the SSH ProxyCommand, as shown in the following line:

```
ssh -i /<path>/<ssh_key> -o ProxyCommand='ssh -i /<path>/<ssh_key> -W %h:%p -p 443...' Press Enter again to activate the console.
```

3- Troubleshooting Instances from Instance Console Connections To boot into maintenance mode

Reboot the instance from the Console.

When the reboot process starts, switch back to the terminal window, and you see Console messages start to appear in the window. As soon as you see the GRUB boot menu appear, use the up/down arrow

key to stop the automatic boot process, enabling you to use the boot menu.

In the boot menu, highlight the top item in the menu, and type `e` to edit the boot entry.

In edit mode, use the down arrow key to scroll down through the entries until you reach the line that starts with either `linuxefi` for instances running Oracle

Autonomous Linux 7.x or Oracle Linux 7.x,



or kernel for instances running Oracle Linux 6.x. At the end of that line, add the following: init=/bin/bash  
Reboot the instance from the terminal window by entering the keyboard shortcut CTRL+X.

**NEW QUESTION 176**

You are a solutions architect for a global health care company which has numerous data centers around the globe. Due to the ever growing data that your company is storing, you were Instructed to set up a durable, cost effective solution to archive you data from your existing on-premises tape based backup Infrastructure to Oracle Cloud Infrastructure (OCI).

What is the most-effective mechanism to Implement this requirement?

- A. Use the File Storage Service in OCI and copy the data from your existing tape based backup to the shared file system
- B. Setup an on premises OCI Storage Gateway which will back up your data to OCI Object Storage Archive tier.(Correct)
- C. Setup an on premises OCI Storage Gateway which will back up your data to OCI object Storage Standard tie
- D. Use Object Storage life cycle policy management to move any data older than 30 days from Standard to Archive tier.
- E. Setup an on-promises OCI Storage Gateway which will back up your data to OCI Object Storage Standard
- F. Setup fastConnect to connect your on premises network to your OCI VCN and use rsync tool to copy your data to OCI Object Storage Archive tier.

**Answer:** B

**Explanation:**

Oracle Cloud Infrastructure offers two distinct storage tiers for you to store your unstructured data. Use the Object Storage Standard tier for data to which you need fast, immediate, and frequent access. Use the Archive Storage service's Archive tier for data that you access infrequently, but which must be preserved for long periods of time. Both storage tiers use the same manageable resources (for example, objects and buckets). The difference is that when you upload a file to Archive Storage, the object is immediately archived. Before you can access an archived object, you must first restore the object to the Standard tier. you can use Storage Gateway to move files to Oracle Cloud Infrastructure Archive Storage as a cost effective backup solution. You can move individual files and compressed or uncompressed ZIP or TAR archives. Storing secondary copies of data is an ideal use case for Storage Gateway.

**NEW QUESTION 178**

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.

**Answer:** BCE

**Explanation:**

<https://docs.cloud.oracle.com/en-us/iaas/Content/Database/Tasks/mig-onprembbackup.htm>

**NEW QUESTION 183**

A cloud consultant is working on a implementation project on Oracle Cloud Infrastructure (OCI). As part of the compliance requirements, the objects placed in OCI 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"] },
 "timeAmount": 5, "timeUnit": "DAYS", "isEnabled": true }
]
```

What will happen after this policy is applied?

- A. All the objects having file extension "doc" will be archived for 5 days and will be deleted 10 days after object creation.
- B. All objects with names starting with "doc" will be deleted after 5 days of 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.

**Answer:** B

**NEW QUESTION 184**

.....

## Thank You for Trying Our Product

\* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

\* One year free update

You can enjoy free update one year. 24x7 online support.

\* Trusted by Millions

We currently serve more than 30,000,000 customers.

\* Shop Securely

All transactions are protected by VeriSign!

**100% Pass Your 1z0-997-20 Exam with Our Prep Materials Via below:**

<https://www.certleader.com/1z0-997-20-dumps.html>