1Z0-997-20 Q&As

Oracle Cloud Infrastructure 2020 Architect Professional

Pass Oracle 1Z0-997-20 Exam with 100% Guarantee

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

https://www.pass4itsure.com/1z0-997-20.html

100% Passing Guarantee
100% Money Back Assurance

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

- **Instant Download** After Purchase
- **100% Money Back** Guarantee
- **365 Days** Free Update
- **800,000+** Satisfied Customers

QUESTION 1

You are working as a solution architect for an online retail store to create a portal to allow the users to pay for their groceries using credit cards. Since the application is not fully compliant with the Payment Card Industry Data Security Standard (PCI DSS), your company is looking to use a third party payment service to process credit card payments.

The third party service allows a maximum of Spelunk IP addresses 5 public IP addresses at a time However, your website is using Oracle Cloud Infrastructure (OCI) Instance Pool Auto Scaling policy to create up to create up to 15 Instances during peak traffic demand, which are launched In VCN private in VCN private subnets and attached to an OCI public Load Balancer. Upon user payment, the portal connects to the payment service over the Interne! to complete the transaction

What solution can you implement to make sure that all compute Instances can connect to the third party system to process the payments aw peak traffic demand?

A. Route credit card payment request from the compute instances through the NAT Gateway. On the third-party services, whitelist the public IP associated with the NAT Gateway.

B. Create an OCI Command Line Interface (CLI) script to automatically reserve public IP address for the compute instances. On the third-party services, whitelist the Reserved public IP.

C. Whitelist the Internet Gateway Public IP on the third party service and route all payment requests through the Internet Gateway.

D. Route payment request from the compute instances through the OCI Load Balancer, which will then be routed to the third party service.

Correct Answer: A

Explanation: https://docs.oracle.com/en-us/iaas/Content/Balance/Concepts/balanceoverview.htm

QUESTION 2

Which three scenarios are suitable for the Oracle Infrastructure (OCI) Autonomous transaction Processing Server less (ATP-S) deployment?

A. well established, online auction marketplace is running an application where there is database usage 24? but also has peaks of activity that the hard to predict when the peaks happen, the total activities may reach 3 times the normal activity level

B. A small startup is deploying a new application fen eCommerce and it requires database to store customers\\' transactions the team b of what the load will look like since it is a new application.

C. A midsize company is considering migrating its legacy on premises MongoDB database to Oracle Cloud Infrastructure (OCI). The database has significantly higher workloads on weekends than weekdays

D. A developer working on an Internal project needs to use a database during work hours but doesn\\'t need It during nights or weekends. the project budget requires her to keep costs low.

E. A manufacturing company is running Oracle E-Business Suite application on premises. They are looking to move this application to OCI and they want to use a managed database offering for their database tier.

Correct Answer: ABD

MongoDB is a cross-platform document-oriented database program. Classified as a NoSQL database program, MongoDB uses JSON-like documents with schema, so the best to be migrated to Oracle NoSQL Database. https://blogs.oracle.com/nosql/migrate-mongodb-data-to-oracle-nosql-database Autonomous transaction Processing Serverless (ATP-S) isn\\'t supported yet for EBS database

QUESTION 3

You have been asked to create a mobile application which will be used for submitting orders by users of a popular E-Commerce site. The application is built to work with Autonomous Transaction Processing - Serverless (ATP-S) database as the backend and HTML5 on Oracle Application Express as the front end. During the peak usage of the application you notice that the application response time is very slow. ATP-S database is deployed with 3 CPU cores and 1 TB of memory.

Which two options are expensive or impractical ways to improve the application response times?

A. Identify the maximum memory capacity needed for peak times and scale the memory for the ATP-S database to that number. ATP-S will scale the memory down when not needed.

B. Use the Machine Learning (ML) feature of the ATP-S database iteratively to tune the SQL queries used by the application.

C. Scale up CPU core count and memory during peak times.

D. Enable auto scaling for CPU cores on ATP-S database.

E. Identify the maximum CPU capacity needed for peak times and scale the CPU core count for the ATP-S database to that number. ATP-S will scale the CPU core count down when not needed.

Correct Answer: CE

QUESTION 4

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

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

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

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

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

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

D. Launch a compute instance and run an NGINX server to host the application. Deploy a 2 node VM DB Systems with Oracle RAC enabled. Setup Oracle GoldenGate to synchronize data from their on-premises database to OCIVM

Database. Export and Import the on- premises database to OCIVM DB Systems using Oracle Data Pump, apply the GoldenGate trail files to sync up the OCI database with the on-premises database. Enable automatic backups for the OCIVM database and then cutoverthe application from on-premises to OCI.

Correct Answer: D

QUESTION 5

A company has an urgent requirement to migrate 300 TB of data to Oracle Cloud Infrastructure (OCI) In two weeks. Their data center has been recently struck by a massive hurricane and the building has been badly damaged, although still operational. They have a 100 Mbps Internet line but the connection is Intermittent due to the damages caused to the electrical grid.

In this scenario, what is the most effective service to use to migrate the data to OCI given the time constraints?

A. Setup a OCI Storage Gateway to connect your data center and your VCN. Once the connection has been established, upload all data to OCI using OCI Storage Gateway Cloud Sync tool.

B. Setup a hybrid network by launching aIGbpsFastConnect virtual circuit between your data center and OCI. Use OCI Object storage multipart upload tool to automate the migration of your data to OCI.

C. Use multiple OCI Data Transfer Appliances to transfer data to OCI.

D. Upload the data to OCI using OCI Object Storage multipart upload tool.

E. Storage Gateway to connect your data center and your VCN. Once the connection has been established, upload all data to OCI.

Correct Answer: C

Due to the network speed is not good enough and the connection is Intermittent due to the damages caused to the electrical grid Oracle offers offline data transfer solutions that let you migrate data to Oracle Cloud Infrastructure. You have 2 Options of Data Transfer DISK-BASED DATA TRANSFER You send your data as files on encrypted commodity disk to an Oracle transfer site. Operators at the Oracle transfer site upload the files into your designated Object Storage bucket in your tenancy. APPLIANCE-BASED DATA TRANSFER you send your data as files on secure, high-capacity, Oracle-supplied storage appliances to an Oracle transfer site. Operators at the Oracle transfer site upload the data into your designated Object Storage bucket in your tenancy.

QUESTION 6

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 Al 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 Al 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 Al 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

Correct Answer: D

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.

QUESTION 7

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 webservers

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

Correct Answer: B

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.

QUESTION 8

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 i his requirement?

A. Object Storage

B. File System Storage

C. Archive storage

D. Block Volume

Correct Answer: B

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

QUESTION 9

A retail company has recently adopted a hybrid architecture. They have the following requirements for their end-to-end Connectivity model between their on-premises data center and Oracle Cloud Infrastructure (OC1) region

*

Highly available connection with service level redundancy

*

Dedicated network bandwidth with low latency

Which connectivity setup is the most cost effective solution for this scenario?

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 connection. Use separate edge devices in your FastConnect physical connectivity is redundant Use a single edge device in your on premises data center for each connection From yc device, advertise more specific routes via primary FastConnect virtual circuit, and less specific routes through t backup FastConnect circuit.

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

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

Correct Answer: D

Explanation: there are two main requirements for this Customer First Highly available connection with service level redundancy and that can achieve by

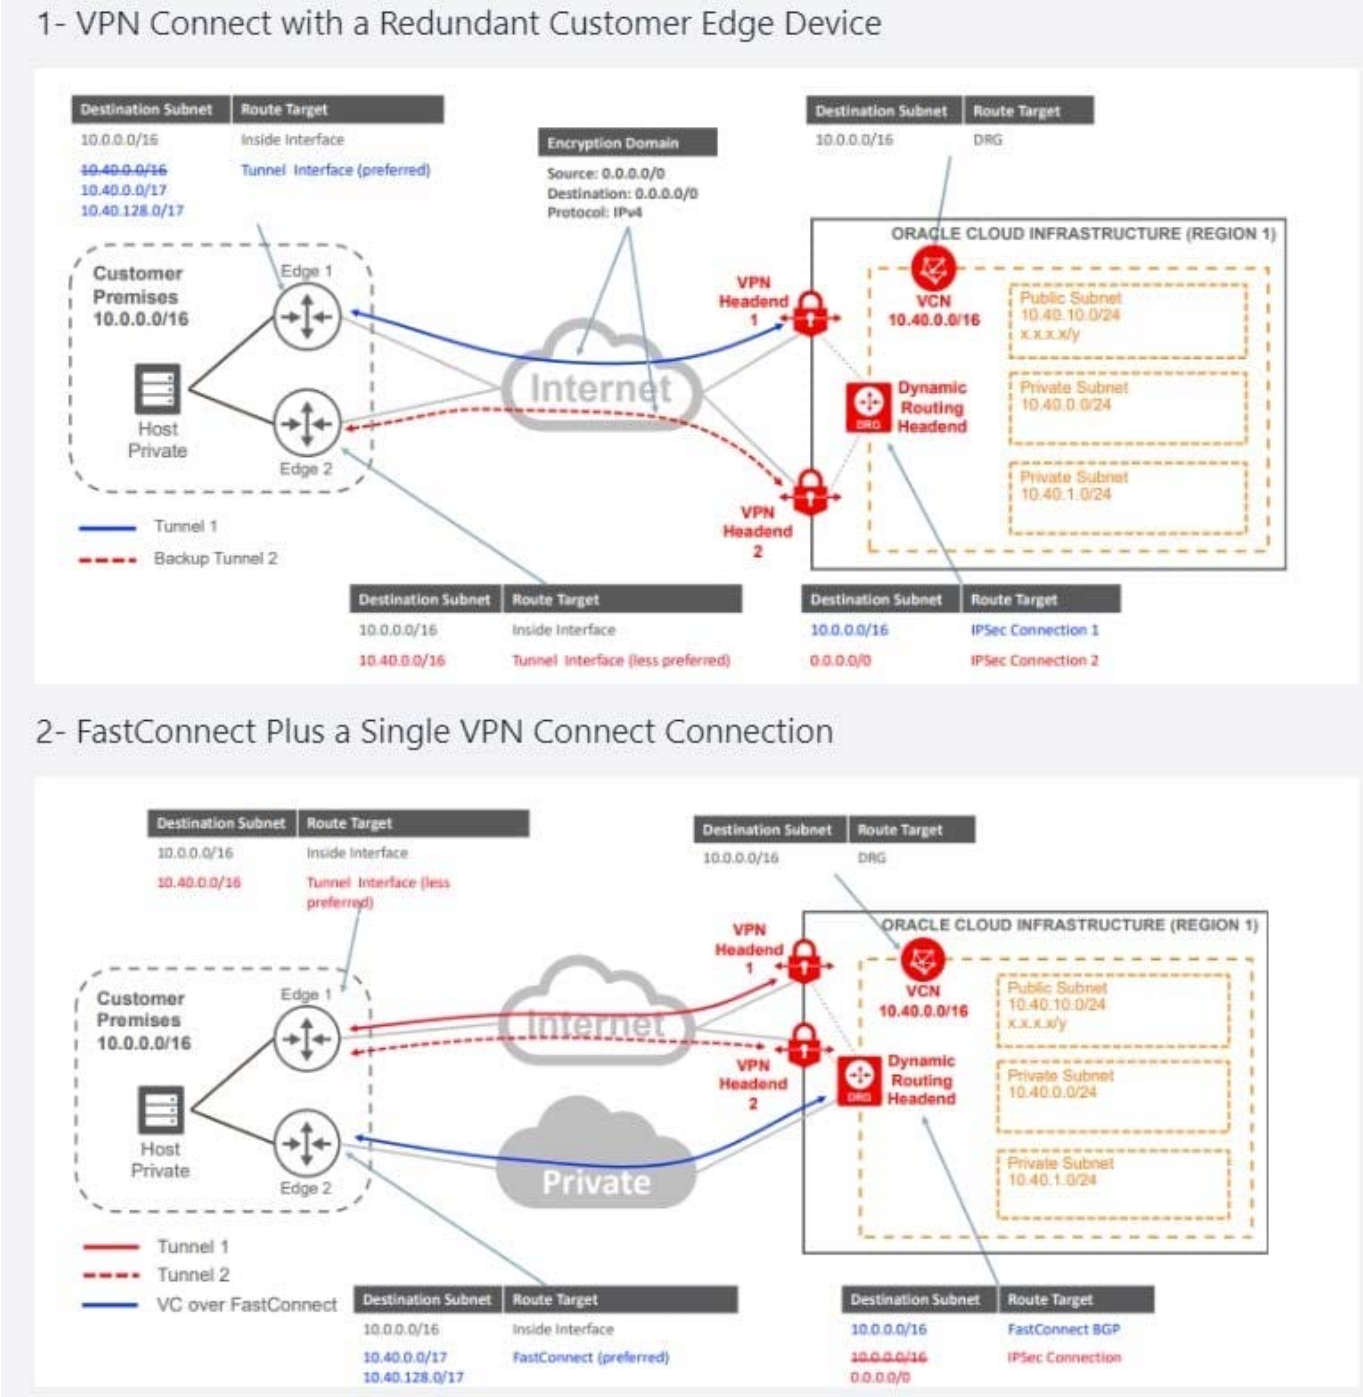

3- Redundant FastConnect

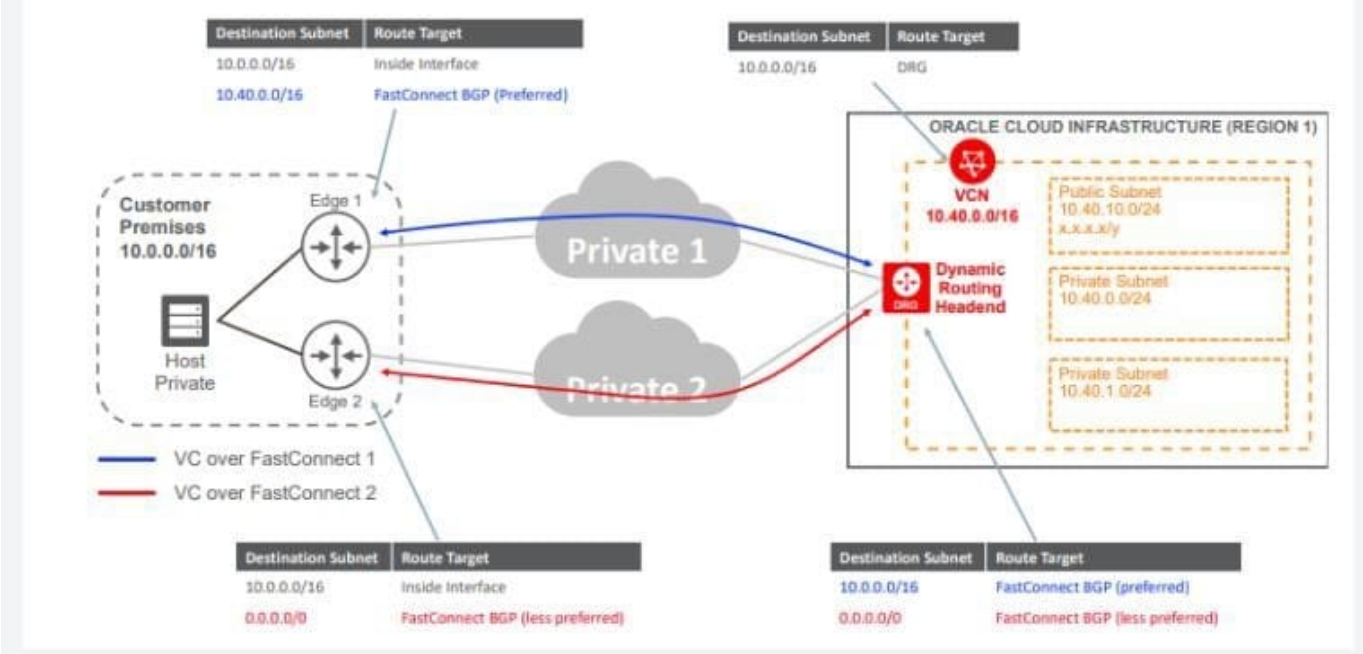

QUESTION 10

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 rules. 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 Gateway. Associate the routing table with all the VCN subnets.

B. 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.

C. 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.

D. 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.

Correct Answer: C

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.

QUESTION 11

Your customer has gone through a recent departmental re structure. As part of this change, they are organizing their Oracle Cloud Infrastructure (OCI) compartment structure to align with the company\\'s new organizational structure.

They have made the following change:

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

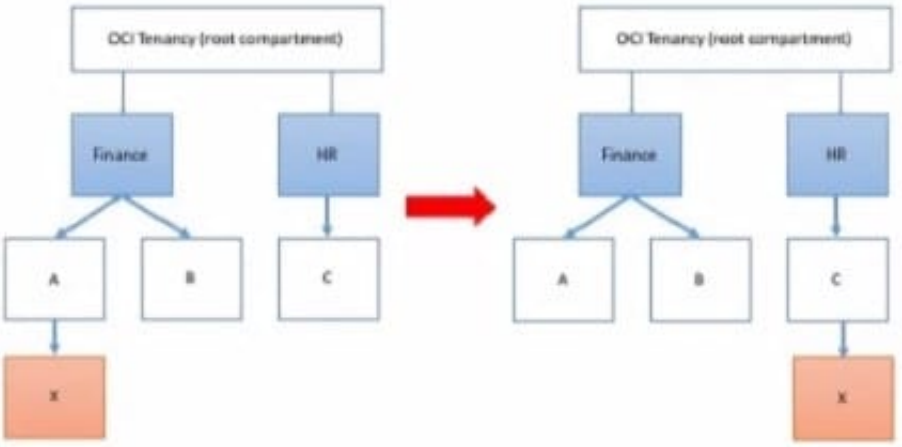

Policy defined in compartment A: Allow group networkadmins to manage subnets in compartment X Policy defined in root compartment: Allow group admins to read subnets in compartment Finance:A:X After you move the compartment, which two IAM policies would be required to ensure both groups retain the same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

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

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

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

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

Correct Answer: BC

QUESTION 12

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 closets 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

Correct Answer: BE

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 ly configured or having other issues. Although the health check appears good customers might experience transaction failures.

QUESTION 13

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

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

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

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

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

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

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

Correct Answer: D

QUESTION 14

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.

Correct Answer: A

QUESTION 15

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. 2. Run kubectl cordon against all the worker nodes in the old pool to stop any new application pods to get scheduled 3. Run kubectl drain """"delete""local""data """"force """"ignore""daemonsets to evict any Pods that are running 4. Delete the old node pool

C. 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

D. 1. Run kubectl cordon against all the worker nodes in the old pool to stop any new application pods to get scheduled 2. Run kubectl drain """"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

Correct Answer: B

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