DUMPSARENA

Oracle Cloud Infrastructure 2022 Architect Professional

Oracle 1z0-997-22

Version Demo

Total Demo Questions: 10

Total Premium Questions: 165

Buy Premium PDF

https://dumpsarena.co

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

QUESTION NO: 1

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

QUESTION NO: 2

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

ANSWER: B E

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

QUESTION NO: 3

You have multiple IAM users who launch different types of compute Instances and block volumes every day. As a result, your Oracle cloud Infrastructure (OCF) 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 resources. This will allow a user to launch compute instances on\y if certain tags were defined.

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

D. Create a default tag for each compartment, which ensure that appropriate tags are applied at resource creation

E. 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.

${} 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

QUESTION NO: 4

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. iam.policy.id - the OCID of the IAM policy to apply to the group.

B. instance.tenancy.id - the OCID of the tenancy where the instance resides.

C. tag...value - the tag namespace and tag key.

D. instance.compartment.id - the OCID of the compartment where the instance resides.

ANSWER: C D

Explanation:

You can define the members of the dynamic group based on the following:

- compartment ID

- instance ID

- tag namespace and tag key

- tag namespace, tag key, and tag value

Supported variables are:

instance.compartment.id - the OCID of the compartment where the instance resides

instance.id - the OCID of the instance

tag...value - the tag namespace and tag key. For

example, tag.department.operations.value .

tag...value='' - the tag namespace, tag key, and tag value. For

example, tag.department.operations.value='45'

QUESTION NO: 5

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-A37A-121241D21199

QUESTION NO: 6

You work for a German company as the Lead Oracle Cloud Infrastructure architect. You have designed a highly scalable architecture for your company's business critical application which uses the Load Balancer service auto which uses the Load Balancer service, autoscaling configuration for the application servers and a 2 Node VM Oracle RAC database. During the peak utilization period of the- application yon notice that the application is running slow and customers are complaining. This is resulting in support tickets being created for API timeouts and negative sentiment from the customer base.

What are two possible reasons for this application slowness?

A. Autoscaling configuration for the application servers didn't happen due to IAM policy that's blocking access to the application server compartment

B. The Load Balancer configuration is not sending traffic to the listener of the application servers.

C. Autoscaling configuration for the application servers didn't happen due to compartment quota breach of the VM shapes used by the application servers.

D. Autoscaling configuration for the application servers didn't happen due to service limit breach of the VM shapes used by the application servers

E. The Load Balancer doesn't have a Network Security Group to allow traffic to the application servers.

ANSWER: C D

Explanation:

Autoscaling

Autoscaling enables you to 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.

Prerequisites

- You have an instance pool. Optionally, you can attach a load balancer to the instance pool. For steps

to create an instance pool and attach a load balancer, see Creating an Instance Pool.

- Monitoring is enabled on the instances in the instance pool. For steps to enable monitoring, see

Enabling Monitoring for Compute Instances.

- The instance pool supports the maximum number of instances that you want to scale to. This limit is

determined by your tenancy's service limits.

About Service Limits and Usage

When you sign up for Oracle Cloud Infrastructure, a set of service limits are configured for your tenancy.

The service limit is the quota or allowance set on a resource. For example, your tenancy is allowed a maximum number of compute instances per availability domain. These limits are generally established with your Oracle sales representative when you purchase Oracle Cloud Infrastructure.

Compartment Quotas

Compartment quotas are similar to service limits; the biggest difference is that service limits are set by Oracle, and compartment quotas are set by administrators, using policies that allow them to allocate resources with a high level of flexibility.

QUESTION NO: 7

An organization has its mission critical application consisting of multiple application servers and databases running inside Virtual Cloud Network (VCN) in uk-london-1 region. Their solution architect wants to further strengthen their architecture by planning for Disaster Recovery (DR) in eu-frankfurt-1 region.

Which two solutions should their architect keep in mind while designing for DR?

A. A remote VCN peering connection is required to establish secure and reliable connectivity between different VCNs created in uk-london-1 and eu-frankfurt-1 region.

B. rsync utility can be used to asynchronously copy file systems or snapshot data to another region.

C. Load balancer will automatically distribute traffic between both the regions.

D. The RTO is the acceptable timeframe of lost data that application can tolerate.

E. It is not possible to use Active Data Guard to synchronize a database in uk-london-1 region to equivalent database in eu-frankfurt-1 region.

ANSWER: A C

QUESTION NO: 8

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 FSS. Add stateless ingress and egress rules for specific P 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.

ANSWER: C

Explanation:

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

QUESTION NO: 9

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

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

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

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

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

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

ANSWER: A

QUESTION NO: 10

You are building a demo for a customer that showcases Oracle Cloud Infrastructure (OCI) Events service and Oracle Functions. You plan to create an event every time an image is uploaded to an OCI Object Storage bucket. You have also created a function that is listening to the event and processes the image for face recognition.

Choose the two actions from below that are NOT required to run the demo successfully.

A. You must specify an action type while creating an Event service and specify the function you want to trigger.

B. Creating an event rule is not permitted for OCI Object storage.

C. The function must be deployed only to Oracle Kubernetes Engine (OKE).

D. You have to enable Object Storage buckets to emit events for state changes.

E. You must deploy the function that does facial recognition for the demo to work.

ANSWER: B C