2022 Realistic ExamsLabs 1z0-997-21 Dumps PDF - 100% Passing Guarantee [Q68-Q90

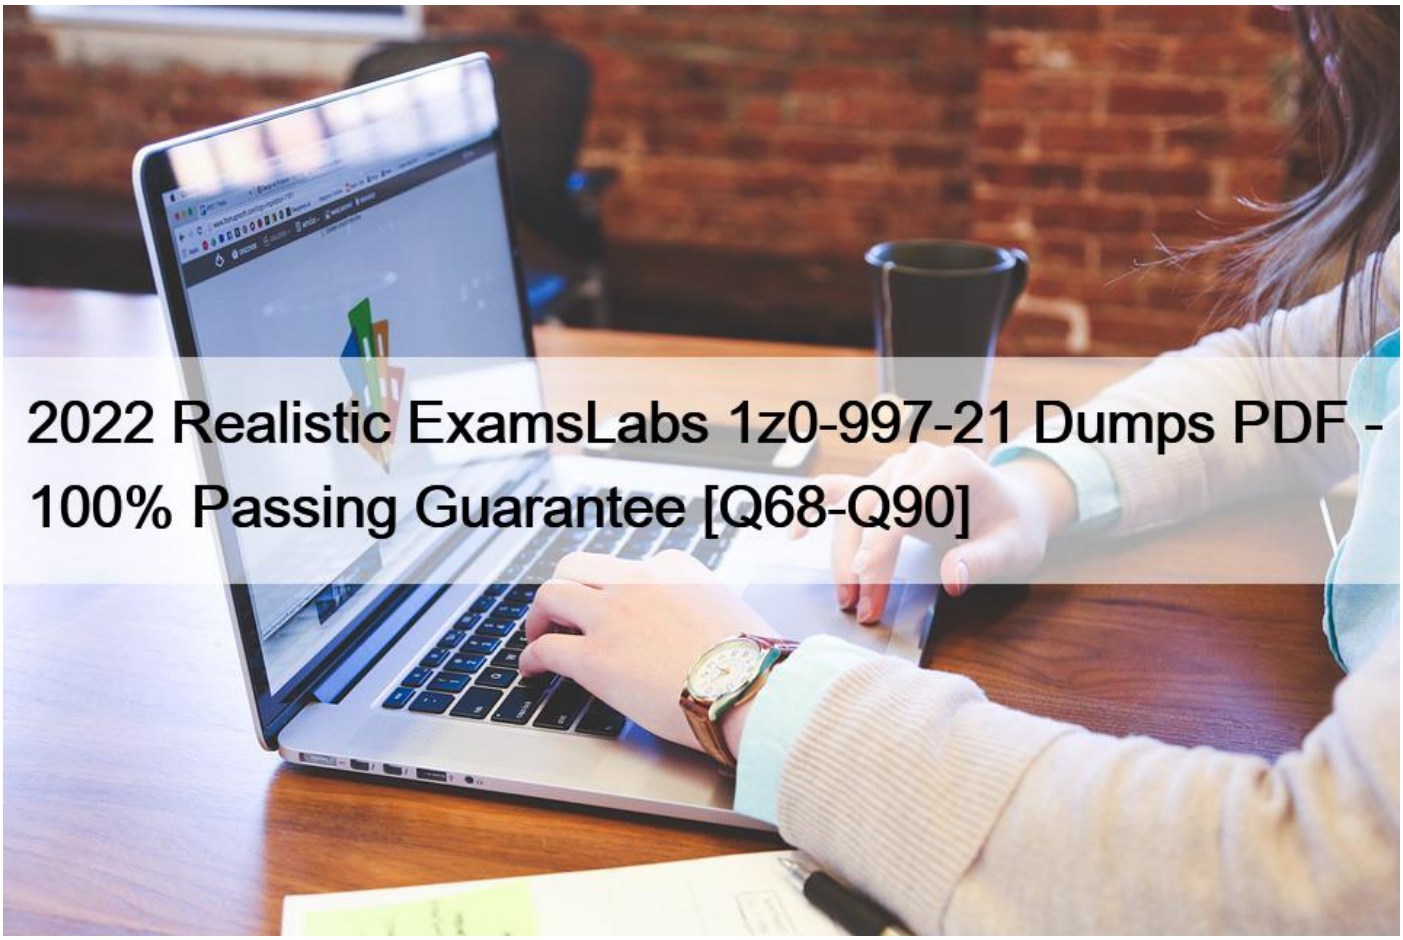

2022 Realistic ExamsLabs 1z0-997-21 Dumps PDF - 100% Passing Guarantee
Free Oracle 1z0-997-21 Exam Questions and Answer

NEW QUESTION 68

A manufacturing company is planning to migrate their on-premises database to Oracle Cloud Infrastructure and has hired you for the migration. Customer has provided following information regarding their existing on-premises database:

Database 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)
* On-Premises host operating system and version.
* Number of active connections.
* Data types used in the on-premises database.
* Elapsed time since database was last patched.
* Top 5 longest running queries.

NEW QUESTION 69

A telecom company has an application running in Oracle Cloud Infrastructure (OCI) Germany Central (eu-frankfurt-1) region. They want to configure Disaster Recovery (DR) site in the OCI UK South (uk-london-1) region. Which is the most cost effective option to help set up application and persistence layers in the DR site?
* Application layer: configure events service rule in eu-frankfurt-1 region to filter Health Checks event failure and route traffic to uk-london-1 region in the event of a disaster.

Persistence layer: set up policy to schedule cross-region automated backups of block volumes between eu-frankfurt-1 and uk-london-1 regions.
* Application layer: configure Traffic Management steering policy with Load Balancing policy between servers in eu-frankfurt-1 and uk-london-1 regions.

Persistence layer: set up policy to schedule cross-region automated backups of block volumes between eu-frankfurt-1 and uk-london-1 regions.
* Application layer: Set us a public laod balancerin the eu-frankfurt-1 region. Create a backend set with instances running in bothuk-frankfurt-1 and uk-london-1 regions.

Persistence layer: Set up OCI Object Storage replication from eu-frankfurt-1 region to uk-london-1 region.
* Application layer: configure Traffic Management steering policy with Failover policy between servers in eu-frankfurt-1 and uk-london-1 regions.

Persistence layer: set up policy to schedule cross-region automated backups of file systems in File Storage service between eu-frankfurt-1 and uk-london-1 regions.

NEW QUESTION 70

You developed a microservices based application that runs on Oracle Cloud Infrastructure (OCI) Container Engine for Kubernetes (OKE). Your security team wants to use SSL termination for this application. What should you do to create a secure SSL termination for this application using fewest steps?
* Create a self-signed certificate and it’s corresponding key. Create a Kubernetes secret using the certificate and the key. Then add these an notations to the Kubernetes service:

annotations:

service.beta.kubernetes.io/oci-load-balancer-ssl-ports: “443”

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

annotations:

service.beta.kubernetes.io/oci-load-balancer-ssl-ports: “443”

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

Service.beta.kubernete.io/oci-load-balancer-ssl-ports: “443”

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

NEW QUESTION 71

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

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

What could be the cause?
* You are exceeding your volume group backup quota configured.
* You have the same block volume attached to multiple compute instances; if these compute instances are in different compartments then all concerned compartments must be moved at the same time.
* Compute instance with multiple block volumes attached cannot move when a compartment is moved.
* The Identity and Access Management policy allowing backup failed to move when the compartment was moved.

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

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

You are aware of the policies that govern access to the compartment in its current position.

You are aware of the polices in the new parent compartment that will take effect when you move the compartment.

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

NEW QUESTION 72

A manufacturing company is planning to migrate their on-premises database to OCI and has hired you for the migration. Customer has provided following information regarding their existing onpremises database:

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

What additional information do you need from customer in order to recommend a suitable migration method? Choose two
* Elapsed time since database was last patched
* On-premises host operating system and version
* Number of active connections
* Data types used in the on-premises database
* Top 5 longest running queries

Not all migration methods apply to all migration scenarios. Many of the migration methods apply only if specific characteristics of the source and destination databases match or are compatible. Moreover, additional factors can affect which method you choose for your migration from among the methods that are technically applicable to your migration scenario.

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

On-premises database version

Database service database version

On-premises host operating system and version

On-premises database character set

Quantity of data, including indexes

Data types used in the on-premises database

Storage for data staging

Acceptable length of system outage

Network bandwidth

NEW QUESTION 73

You have provisioned a new VM.DenseIO2.24 compute instance with local NVMe drives. The compute instance is running production application. This is a write heavy application, with a significant Impact to the business it the application goes down.

What should you do to help maintain write performance and protect against NVMe devices failure.
* NVMe drive have built in capability to recover themself so no other actions are required
* Configure RAID 6 for NVMe devices.
* Configure RAID 1 for NVMe devices.
* Configure RAID 10 for NVMe devices.
VM.DeselO2.24 compute instance include locally attached NVMe devices. These devices provide extremely low latency, high performance block storage that is ideal for big data, OLTP, and any other workload that can benefit from high-performance block storage.

A protected RAID array is the most recommended way to protect against an NVMe device failure. There are three RAID levels that can be used for the majority of workloads:

RAID 1: An exact copy (or mirror) of a set of data on two or more disks; a classic RAID 1 mirrored pair contains two disks RAID 10: Stripes data across multiple mirrored pairs. As long as one disk in each mirrored pair is functional, data can be retrieved RAID 6: Block-level striping with two parity blocks distributed across all member disks If you need the best possible performance and can sacrifice some of your available space, then RAID 10 array is an option.

NEW QUESTION 74

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?
* Create two private virtual circuits on the FastConnect link. Create two Dynamic Routing Gateways, one for each VCNs. Attach the virtual circuits to the dynamic routing gateways.
* You cannot achieve connectivity using single FastConnect link as the production and the development VCNs-are in separate

tenancies. Request one more FastConnect connection.
* Create a single private virtual circuit over FastConnect and attach fastConnect to either of the VCN’s Dynamic Routing Gateway. Use Remote Peering to peer production and development VCNs.
* Create a hub-VCN that uses Dynamic Routing Gateway (DRG) to communicate with on-premises network over FastConnect. Connect the hub-VCN to the production VCN spoke and with development VCN spoke, each peered via their respective local Peering Gateway (LPG)
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 75

A startup company is looking for a solution for processing of data transmitted by the IOT devices fitted to transport vehicles that carry frozen foods. The data should be consumed and processed in real time. The processed data should be archived to OCI Object Storage bucket. and use Autonomous Data warehouse (ADW) to handle analytics.

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

– OCI Streaming service provides a fully managed, scalable, durable storage option for continuous, highvolume streams of data that you can consume and process in real-time

– Use cases

Log and Event data collection

Web/Mobile activity data ingestion

IoT Data streaming for processing and alerts

Messaging: use streaming to decouple components of large systems

– Oracle managed service with REST APIs (Create, Put, Get, Delete)

– Integrated Monitoring

NEW QUESTION 76

You are working as a security consultant with a global insurance organization which is using Microsoft Azure Active Directory (AD) as identity provided to manager user login/passwords. When a user logs in to Oracle Cloud infrastructure (OCI) console, it should get authenticated by Azure AD.

Which set of steps are required to configure at OCI side in order to get it enabled
* Setup Azure AD as an Enterprise Application, map Azure AD users and groups and policies to OCI groups and users
* Setup Azure AD as an Identity Provider, Import users and groups from Azure AD to OCI, set up IAM policies to govern access to Azure AD groups
* Setup Azure AD as an Enterprise Application, configure OCI for single sign-on, map Azure AD groups to OCI groups, set up the IAM policies to govern access to Azure AD groups
* Setup Azure AD as an Identity Provider, map Azure AD groups to OCI groups, set up the IAM policies to govern access to Azure AD groups
Federating with Microsoft Azure Active Directory

To federate with Azure AD, you set up Oracle Cloud Infrastructure as a basic SAML single sign-on application in Azure AD. To set up this application, you perform some steps in the Oracle Cloud Infrastructure Console and some steps in Azure AD.

Following is the general process an administrator goes through to set up the federation. Details for each step are given in the next section.

In Oracle Cloud Infrastructure, download the federation metadata document.

In Azure AD, set up Oracle Cloud Infrastructure Console as an enterprise application.

In Azure AD, configure the Oracle Cloud Infrastructure enterprise application for single sign-on.

In Azure AD, set up the user attributes and claims.

In Azure AD, download the Azure AD SAML metadata document.

In Azure AD, assign user groups to the application.

In Oracle Cloud Infrastructure, set up Azure AD as an identity provider.

In Oracle Cloud Infrastructure, map your Azure AD groups to Oracle Cloud Infrastructure groups.

In Oracle Cloud Infrastructure, set up the IAM policies to govern access for your Azure AD groups.

Share the Oracle Cloud Infrastructure sign-in URL with your user

NEW QUESTION 77

As a solution architect, you are designing a web application to be deployed across multiple Oracle Cloud Infrastructures (OCI) regions for a global audience. Your goal is that users from each region should access the application web servers deployed in their own geographical OCI location.

Which OCI feature can be used to achieve this?
* OCI Traffic Management IP Prefix steering policy
* OCI Global Load balancers
* OCI Public Load Balancers
* OCI Traffic Management GeoLocation steering policy

NEW QUESTION 78

A cost conscious fashions design company which sells bags, clothes, and other luxury items has recently decided to more all of the their on-premises infrastructure 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.

Based on cost considerations, what is an effective mechanism to migrate the customer application to OCI and set up regular automated backups?
* Launch a compute Instance and run a NGINX server to host the application. Deploy a 2 node VM DB Systems with oracle RAC enabled import the on premises database to OCI VM DB Systems using oracle Data Pump and then enable automatic backups.
* Launch a compute Instance and run an NGINX server to host the application. Deploy Exadata Quarter Rack, enable automatic backups and import the database using Oracle Data Pump.
* Launch a compute Instance for both the NGINX application server and the database server. Attach block volumes on the database server compute instance and enable backup policy to backup the block volumes.
* Launch a Compute instance and run a 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 data pump and then enable automatic backup- Also, enable Oracle Data Guard on the database server
Based on cost considerations will exclude the Exadata. and there’s no need for Data Guard Cost Estimator

https://www.oracle.com/cloud/cost-estimator.html

NEW QUESTION 79

An insurance company is storing critical financial data in the Oracle Cloud Infrastructure block volume. This volume is currently encrypted using oracle managed keys. Due to regulatory compliance, the customer wants to encrypt the data using the keys that they can control and not the keys which are controlled by Oracle.

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

NEW QUESTION 80

You are using the Oracle Cloud Infrastructure (OCI) OS Management service to manage updates and patches for the Oracle Linux 8 environments on your compute instances in OCI. You have verified that the OS Management Service Agent (osms-agent) is installed and running properly in the instances.

One of the compute instances is not getting the updates from OS Management Service. You use the following command to validatethatyourinstance cannot reach the OS Management Ingestion service by running curl https://ingestion.osms.

<region>.oci.oraclecloud.com/

Which Is NOT a possible reason for this issue?
* The instance Is in a private subnet with a NAT gateway.
* The instance is in a private subnet with a privateendpointwith security rules configured to access the OS Managementingestion service
* The instance is in a private subnet with a service gatewaythat uses the All <region> Services in Oracle Services Network CIDR label.
* The Instance is in a publicsubnet with an Internet gateway.

NEW QUESTION 81

You have created compartment called Dev for developers. There are two IAM groups for developers: group-devl and group-dev2. You need to write an Identity and Access Management (IAM) policy to give users in these groups access to manage all resources in the compartment Dev.

Which of the following IAM policy will accomplish this?
* Allow any-user to manage all resources in compartment Dev where request.group= /group-dev*/
* Allow group group-devl group-dev2 to manage all resources in compartment Dev
* Allow group /group-dev*/ to manage all resources in compartment Dev
* Allow any-user to manage all resources in tenancy where target.comparment= Dev

NEW QUESTION 82

Your company developed a function that needs to access the Oracle Database to inject some data to it at runtime. You are tasked to move this function to the Oracle Cloud Infrastructure (OCI) and use Oracle Functions and access Oracle Autonomous Database. You created a Dockerfile below to run this function, however, you are getting this error “cx_Oracle.DatabaseError:

ORA””12560: TNS:protocol adapter error”.

Dockerfile:

```
FROM oraclelinux:7-slim

RUN yum -y install oracle-release-el7 oracle-nodejs-release-el7 && \
yum-config-manager --disable ol7_developer_EPEL && \
yum -y install oracle-instantclient19.3-basiclite nodejs && \
rm -rf  /var/cache/yum

WORKDIR /function
ADD . /function/
RUN npm install

CMD exec node func.js
```


What should you do to make sure that Oracle Functions can run this Dockerfile properly? (Choose the best answer.)
* Add these two lines to your Dockerfile: groupadd “”””gid 1000 fn && adduser “”””uid 1000 “”””gid fn fn
* Use “”””privileged flag while running the Docker container to add runtime privilege
* Use “”””cap””add=ALL flag while running the Docker container to add runtime capability
* You ned to run this Container as root, so add this line: USER root
https://docs.cloud.oracle.com/en-us/iaas/Content/Functions/Tasks/functionsrunningasunprivileged.htm

NEW QUESTION 83

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.)
* Take a full database backup of the production Autonomous database and create the non-production database from it.
* Create a metadata clone of the production Autonomous Database and create the non-production database from it.
* Create a full clone of the production Autonomous Database and create the non-production database from it.
* Take a Data Pump export of the production Autonomous database and import into the non-production database.
https://www.oracle.com/database/technologies/datawarehouse-bigdata/adb-faqs.html

NEW QUESTION 84

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 servers technologies to run these application instances.

Which option should you recommend to build this new platform?
* Install a kubernetes cluster on OCI and use OCI event service.
* Use Oracle Container Engine for kubernetes, OCI Registry and OCI Functions.
* Use OCI Resource Manager to automate compute Instances provisioning and use OCI Streaming service.
* Use OCI functions, OCI object storage and OCI event service.

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.

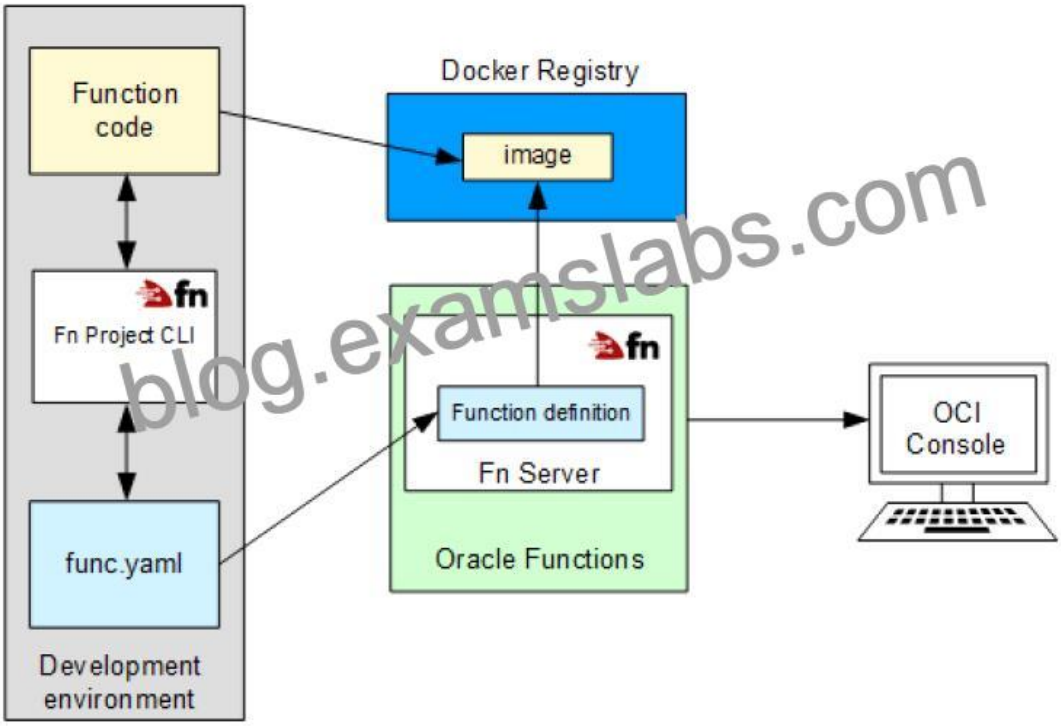

Oracle Cloud Infrastructure Container Engine for Kubernetes is a fully-managed, scalable, and highly available service that you can use to deploy your containerized applications to the cloud. Use Container Engine for Kubernetes (sometimes abbreviated to just OKE) when your development team wants to reliably 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 85

You are working on the migration of the web application infrastructure of your company from on-premises to Oracle Cloud Infrastructure. You need to ensure that the DNS cache entries of external clients will not direct them to the on-premises infrastructure after switching to the new infrastructure.

Which of the following options will minimize this problem?
* Reduce the TTL of the DNS records after the switch.
* DNS changes propagate fast enough that it is not necessary to take any action.
* Increase the TTL of the DNS records before the switch.
* Increase the TTL of the DNS records after the switch.
* Reduce the TTL of the DNS records before the switch.

NEW QUESTION 86

A company has an application that processes confidential dat

a. The data is currently stored in an on-premises data center. A solution architect needs to move this data to Oracle Cloud Infrastructure (OCI) Object Storage and ensure data is encrypted in-transit to OCI.

Which two steps should the solution architect perform to set up the most cost-effective connection between on-premises data center and OCI?
* Set up private end point for accessing Object Storage.
* Attach an Internet Gateway to Virtual Cloud network(VCN).
* Configure a service gateway accessing Object Storage.
* Set up an IPsec tunnel between the customer equipment and software VPN on an oci instance
* Configure a private peering connection on the Oracle FastConnect
* Set up VPN Connect between the customer equipment and the Dynamic Routing Gateway.

NEW QUESTION 87

You are working as a cloud engineer for an IoT startup company which is developing a health monitoring pet collar for dogs and cats. The company collects biometric Information of the pet every second and then sends it to Oracle Cloud Infrastructure (OCI) Your task is to come up with an architecture which will accept and process the monitoring data as well as provide complete trends and health reports to the pet owners. The portal should be highly available, durable, and scalable with an additional feature for showing real time biometric data analytics.

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

NEW QUESTION 88

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?
* 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.
* Use the Machine Learning (ML) feature of the ATP-S database iteratively to tune the SQL queries used by the application.
* Scale up CPU core count and memory during peak times.
* Enable auto scaling for CPU cores on ATP-S database.
* 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.

NEW QUESTION 89

An upcoming e-commerce company has deployed their online shopping application on OCI. 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?
* The health check on some of the backend servers has failed and the load balancer has taken those servers temporarily out of rotation
* 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
* The health check on some of the backend servers has failed and the load balancer was rebooting these servers.
* The autoscaling has already scaled to the maximum number of instances specified in the configuration and there is no room of scaling

NEW QUESTION 90

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.
* 122 for the VCN; 124 for the subnets
* /23 for the VCN; /25 for the subnets
* /24 for the VCN; /24 for the subnets
* /22 for the VCN; /25 for the subnets

Oracle 1z0-997-21 Exam Syllabus Topics:

TopicDetailsTopic 1- Design, implement and operate solutions to meet compliance requirements- Plan and design solutions to meet business and technical requirementsTopic 2- Design, implement and operate databases in OCI- Migrate on-premises workloads to OCI- Operate and troubleshoot solutions on OCITopic 3- Manage infrastructure using OCI CLI, APIs and SDKs- Plan and design solutions in Oracle Cloud Infrastructure (OCI)Topic 4- Evaluate multi-cloud solution architectures- Conduct Monitoring, observability and alerting in OCITopic 5- Create architecture patterns including N-tier applications, microservices, and serverless architectures- Implement and troubleshoot database migrationsTopic 6- Design for Security and Compliance- Evaluate and implement databases- Operate and troubleshoot databases