Oracle.1z0-1084-26.v2026-10-26.q78

Exam Code:	1z0-1084-26
Exam Name:	Oracle Cloud Infrastructure Developer Professional
Certification Provider:	Oracle
Free Question Number:	78
Version:	v2026-10-26
# of views:	111
# of Questions views:	933
https://www.freecram.net/torrent/Oracle.1z0-1084-26.v2026-10-26.q78.html	

NEW QUESTION: 1
Which open source engine is used by Oracle Cloud Infrastructure (OCI) to power Oracle Functions?
A. Knative
B. Kubeless
C. Apache OpenWhisk
D. Fn Project
Answer: (SHOW ANSWER)
Fn Project is the open source engine that is used by OCI to power Oracle Functions1. Fn Project is an open source, container native, serverless platform that can be run anywhere - any cloud or on-premises1. Fn Project is easy to use, extensible, and performant. You can download and install the open source distribution of Fn Project, develop and test a function locally, and then use the same tooling to deploy that function to Oracle Functions1.Verified References: Overview of Functions

NEW QUESTION: 2
As a Cloud Native developer, you develop two services in Node.js and deploy them to two different Container Engine for Kubernetes (OKE) clusters that use the same Virtual Cloud Network (VCN). Your security team wants to analyze the network communication between them. How can this requirement be met in the most cost-effective way?
A. Rewrite the application and send the application logs to an outside log aggregator.
B. Deploy a third-party logging service and aggregate the network flow logs.
C. Use the OCI Logging service and enable VCN flow logs.
D. Deploy Wireshark and intercept the packets.
Answer: (SHOW ANSWER)
The best answer is: " Use the OCI Logging service and enable VCN flow logs. " To meet the requirement of analyzing network communication between two services deployed in different

Container Engine for Kubernetes (OKE) clusters within the same Virtual Cloud Network (VCN) in a cost-effective way, you can use the OCI Logging service and enable VCN flow logs. The VCN flow logs feature in OCI allows you to capture and log network traffic information for your VCN resources. By enabling VCN flow logs, you can monitor and analyze the network communication between your services without the need for additional third- party logging services or tools. Enabling VCN flow logs provides visibility into the network traffic, including source and destination IP addresses, ports, protocols, and other relevant details. This information can be collected and stored in the OCI Logging service, where you can analyze and gain insights into the network communication patterns between your services. By leveraging the built-in capabilities of the OCI Logging service and enabling VCN flow logs, you can fulfill the security team ' s requirement for network communication analysis in a cost-effective manner. This eliminates the need for deploying additional third- party logging services or tools, reducing complexity and potential costs associated with their setup and maintenance. The other options mentioned are not the most cost-effective or suitable solutions for analyzing network communication in this scenario: Deploying a third-party logging service and aggregating the network flow logs would introduce additional costs and complexity, which may not be necessary considering the built- in capabilities provided by OCI. Rewriting the application to send logs to an outside log aggregator would not directly address the requirement of analyzing network communication between the services. It would focus more on application-level logs rather than network-level analysis. Deploying Wireshark and intercepting packets would require additional infrastructure setup and maintenance, which may not be the most cost- effective approach for network analysis in this scenario.

NEW QUESTION: 3
Which of the following step is NOT required for setting up the Container Engine for Kubernetes (OKE) cluster access using a local installation of kubectl?
A. Generate Auth token from the OCI console to access the OKE cluster using kubectl.
B. Install and configure the Oracle Cloud Infrastructure (OCI) CLI.
C. Set up the kubeconfig file.
D. Generate an API signing key pair (if you do not already have one) and upload the public key of the API signing key pair.
Answer: (SHOW ANSWER)
The step that is NOT required for setting up the Container Engine for Kubernetes (OKE) cluster access using a local installation of kubectl is to generate an Auth token from the OCI console. The authentication for accessing the OKE cluster using kubectl can be performed using the OCI CLI configuration, specifically the API signing key pair and the kubeconfig file. Here are the correct steps for setting up the OKE cluster access using a local installation of kubectl: Set up the kubeconfig file: The kubeconfig file contains the necessary information to authenticate and access the OKE cluster using kubectl. It includes details such as the cluster endpoint, authentication method, and credentials. Generate an API signing key pair (if you do not already have one) and upload the public key of the API signing key pair: The API signing key pair is used for authentication with the OCI services. The public key of the key pair needs to be uploaded to the

OCI Console to associate it with your user account. Install and configure the Oracle Cloud Infrastructure (OCI) CLI: The OCI CLI provides a command-line interface to interact with the OCI services. It needs to be installed and configured with your OCI credentials, including the user ' s OCID, tenancy OCID, region, and the path to the API signing key pair. By completing these steps, you can configure kubectl to access and manage your OKE clusters from your local machine using the OCI CLI authentication configuration.

NEW QUESTION: 4
Which statement is NOT valid regarding OCI Identity and Access Management (IAM) policies written for OKE Workload Identity?
A. The policy subject must be specified as any-user rather than a specific IAM group or dynamic group name.
B. The policy condition must explicitly define request.principal.type with the value set to workload.
C. The policy statement must be scoped strictly at the tenancy level and cannot be written to restrict access to a specific compartment.
D. The policy condition must specify the Kubernetes namespace, service account name, and OKE cluster OCID.
Answer: (SHOW ANSWER)
Oracle ' s documented OKE Workload Identity policy pattern is Allow any-user to < verb > < resource > in < location > where all {...} with conditions identifying request.principal.type= ' workload ' , the Kubernetes namespace, service account, and cluster OCID. Crucially, Oracle defines < location > as either the entire tenancyor a specific compartment. Oracle even provides examples granting a workload access only to resources in a named compartment. Consequently, C is false because workload-identity policies do not have to be tenancy-wide. Options A, B, and D accurately describe the standard policy structure: any-user is narrowed through workload-specific conditions, and those conditions identify the exact OKE workload by cluster, namespace, and service account. This enables fine-grained least-privilege access without storing OCI credentials inside pods.

NEW QUESTION: 5
Which TWO statements are correct regarding Docker images and containers? (Choose two.)
A. Only one container can be spawned from a given image at a time.
B. An image is a collection of immutable layers whereas a container is a running instance of an image.
C. If multiple containers are spawned from the same image, then they all use the same copy of that image in memory.
D. Writing and building a new Dockerfile is the only way you can create new Docker images.
E. A container can exist without an image but an image cannot exist without a container.
Answer: (SHOW ANSWER)
Which TWO statements are correct regarding Docker images and containers? (Choose two.)
Writing and building a new Dockerfile is the only way you can create new Docker images. A

container can exist without an image but an image cannot exist without a container. Only one container can be spawned from a given image at a time. If multiple containers are spawned from the same image, then they all use the same copy of that image in memory. (Correct) An image is a collection of immutable layers whereas a container is a running instance of an image. (Correct) ChatGPT The correct statements regarding Docker images and containers are: If multiple containers are spawned from the same image, then they all use the same copy of that image in memory. This statement is correct. When multiple containers are created from the same Docker image, they share the underlying layers of the image in memory. Each container has its own isolated filesystem and runtime environment, but they all use the same base image, which helps in efficient resource utilization. An image is a collection of immutable layers whereas a container is a running instance of an image. This statement is also correct. Docker images are composed of multiple layers, with each layer representing a specific change or addition to the previous layer. These layers are read-only and can be shared among multiple containers. On the other hand, a container is a lightweight and isolated runtime instance created from a specific image. It is a running process that has its own filesystem, network, and runtime configuration. The statements " Writing and building a new Dockerfile is the only way you can create new Docker images " and " A container can exist without an image but an image cannot exist without a container " are incorrect. Docker images can be created in various ways, including using Dockerfiles, importing from existing images, or pulling from registries. Additionally, a container requires an image to run, so an image must exist before a container can be created from it.

NEW QUESTION: 6
Which node configuration in Oracle Cloud Infrastructure Container Engine for Kubernetes provides a serverless experience with granular pod-level elasticity and billing?
A. Managed Nodes
B. Bare Metal Nodes
C. Virtual Nodes
D. Self-Managed Nodes
Answer: (SHOW ANSWER)
OKE Virtual Nodes provide Oracle ' s serverless Kubernetes execution model. Oracle manages the underlying node infrastructure, eliminating the need for customers to provision, patch, upgrade, scale, or troubleshoot conventional worker-node infrastructure. Resource allocation and elasticity occur at the pod level, allowing workloads to scale without planning the capacity of a conventional node pool. Oracle also describes Virtual Nodes as providing pay-per-use pricing based on the compute resources consumed by each Kubernetes pod rather than charging for entire worker servers that might contain unused capacity. Managed, bare metal, and self-managed nodes remain infrastructure-oriented worker-node models and therefore do not provide the same serverless pod-level operating model. Consequently, Virtual Nodes are the configuration designed for granular elasticity and pod-oriented consumption.

NEW QUESTION: 7

You have built a local container image named order-processor:1.0 and need to push it to a repository named order-services in your Phoenix region (phx) tenancy, which has the Object Storage namespace ax39m19vpw.
Which command must you run to properly tag the local image before pushing it to Container Registry?
A. docker tag order-processor:1.0 phx.ocir.io/ax39m19vpw/order-services:1.0
B. docker tag order-processor:1.0 ax39m19vpw/phx.ocir.io/order-services:1.0
C. docker tag order-processor:1.0 ocir.phx.oraclecloud.com/ax39m19vpw/order-services:1.0
D. docker tag order-processor:1.0 ax39m19vpw/order-services/phx.ocir.io:1.0
Answer: (SHOW ANSWER)
OCI Container Registry requires the destination image tag to use the structure < region-key > .ocir.io/ < tenancy-namespace > / < repository-name > : < tag > . For Phoenix, the region key is phx; the supplied tenancy namespace is ax39m19vpw; the repository is order-services; and the image tag is 1.0. Combining these values produces phx.ocir.io/ax39m19vpw/order-services:1.0. Oracle ' s current container documentation provides the same docker tag structure before an image is pushed to OCIR. The alternative commands incorrectly reorder the namespace and registry host or use an unsupported registry hostname format.
Therefore, option A correctly retags the local order-processor:1.0 image with the fully qualified OCIR repository path required for the subsequent docker push.

NEW QUESTION: 8
You are developing a polyglot serverless application using Oracle Functions. Which language cannot be used to write your function code?
A. PL/SQL
B. Python
C. Node.js
D. Go
E. Java
Answer: (SHOW ANSWER)
Oracle Functions does not currently support PL/SQL as a language for writing function code. PL/SQL is a procedural language used in Oracle Database for developing stored procedures, triggers, and other database- related code. However, Oracle Functions supports several other popular programming languages such as Go, Node.js, Python, and Java, allowing developers to choose the language that best suits their application requirements and their familiarity with the language. While PL/SQL is powerful for working with the Oracle Database, it is not an option for writing function code in the Oracle Functions serverless architecture.

NEW QUESTION: 9
Which key pillar of cloud native development should be utilized when a development team needs to automate build, test, and release cycles to achieve rapid and reliable software delivery?
A. Containerization

B. DevOps Principles
C. Microservices Architecture
D. Serverless Computing
Answer: (SHOW ANSWER)
DevOps principles directly address automation across the software delivery lifecycle. Oracle describes DevOps as collaboration between development and operations teams aimed at continuous delivery, with automation applied throughout planning, coding, building, testing, releasing, deploying, operating, and monitoring. OCI specifically identifies CI/CD tooling for automating build, test, release, and deployment activities. Containerization provides consistent application packaging, microservices define an architectural decomposition model, and serverless computing removes much of the underlying infrastructure-management responsibility. None of those technologies by itself defines the automated delivery lifecycle requested in the question. Continuous Integration and Continuous Delivery are core DevOps practices that make software changes repeatable, testable, and deployable at high frequency while reducing manual intervention and release risk. Therefore, the correct cloud-native pillar is DevOps Principles.

NEW QUESTION: 10
You are building an event-driven workflow and need to programmatically create an Oracle Cloud Infrastructure (OCI) Events rule using the OCI Command Line Interface (CLI). Which CLI command correctly creates the rule while sourcing its target destination services from a local configuration file named actions.json?
Freecram.net
A. oci events rule create --display-name " MyRule " --is-enabled true --condition ' { " eventType " : " com.
oraclecloud.objectstorage.deletebucket " } ' --compartment-id ocid1.compartment.oc1..abc --destinations file://actions.json
B. oci events rule create --display-name " MyRule " --is-enabled true --condition ' { " eventType " : " com.
oraclecloud.objectstorage.deletebucket " } ' --compartment-id ocid1.compartment.oc1..abc --targets file://actions.json
C. oci events rule create --display-name " MyRule " --is-enabled true --condition ' { " eventType " : " com.
oraclecloud.objectstorage.deletebucket " } ' --compartment-id ocid1.compartment.oc1..abc --action- payload file://actions.json
D. oci events rule create --display-name " MyRule " --is-enabled true --condition ' { " eventType " : " com.
oraclecloud.objectstorage.deletebucket " } ' --compartment-id ocid1.compartment.oc1..abc --actions file://actions.json
Answer: (SHOW ANSWER)
The OCI CLI command oci events rule create requires the --actions parameter to define the destinations invoked when the Events rule condition matches. Oracle identifies --actions as a required complex parameter and explicitly supports loading its JSON representation from a local

file using file://path/to/file. Oracle ' s Events examples show the same pattern: --actions file://actions.json. The other proposed parameters--- destinations, --targets, and --action-payload- are not valid parameters for the events rule create CLI operation. The actions file can define supported destinations such as Notifications, Streaming, or Functions.
Consequently, option D precisely matches the current OCI CLI command-reference syntax for programmatically creating the rule with action definitions stored in a file.

NEW QUESTION: 11
Who is responsible for patching, upgrading, and maintaining the worker nodes in Oracle Cloud Infrastructure (OCI) Container Engine for Kubernetes (OKE)? (Choose the best answer.)
A. Oracle Support
B. It is automated
C. The user
D. Independent Software Vendors
Answer: C (LEAVE A REPLY)
The user is responsible for patching, upgrading, and maintaining the worker nodes in Oracle Cloud Infrastructure (OCI) Container Engine for Kubernetes (OKE). In OKE, the user has control over the worker nodes, which are the compute instances that run the Kubernetes worker components. As the user, you are responsible for managing and maintaining these worker nodes, including tasks such as patching the underlying operating system, upgrading Kubernetes versions, and performing any necessary maintenance activities. While Oracle provides the underlying infrastructure and support services, including managing the control plane and ensuring the availability of the OKE service, the responsibility for managing the worker nodes lies with the user. This allows you to have control and flexibility in managing your Kubernetes environment according to your specific needs and requirements.

NEW QUESTION: 12
You are developing a cloud native serverless application where uploaded PDF documents inside an OCI Object Storage bucket must automatically trigger an OCI Function for text processing. You have created an OCI Events rule configured with the condition eventType: com.oraclecloud.objectstorage.createbucket matching the bucket ' s compartment. However, when test files are uploaded, the function does not trigger.
Which TWO actions must you perform to successfully implement and troubleshoot this event-driven integration?
A. Define a dynamic group for the Object Storage bucket to allow it to invoke the targeted Oracle Function.
B. Enable the " Emit Object Events " setting on the source Object Storage bucket.
C. Create an IAM policy allowing the OCI Events service to read objects in the specific Object Storage bucket.
D. Verify the rule ' s condition is configured with the com.oraclecloud.objectstorage.createobject event type rather than com.oraclecloud.objectstorage.createbucket.

E. Set up a VCN Service Gateway because OCI Events cannot access Object Storage buckets over the public internet.

Answer: (SHOW ANSWER)

Object-level state changes in OCI Object Storage donot emit Events by default. Oracle requires the bucket ' sEmit Object Eventscapability to be enabled before object create, update, or delete events are generated. In addition, the rule must use the correct event type. com.oraclecloud.objectstorage.createbucket represents creation of a bucket itself, whereas uploading a new object generates com.oraclecloud.objectstorage.
createobject. Therefore, the current rule can never match ordinary PDF uploads. Dynamic groups are not used to assign identities to buckets, Events does not require permission to read the actual object ' s contents, and no VCN Service Gateway is necessary for this native OCI service integration. Consequently, B and D are required.

NEW QUESTION: 13

Which statement is NOT valid regarding autoscaling configurations in Oracle Cloud Infrastructure Container Engine for Kubernetes (OKE)?

A. The Kubernetes Horizontal Pod Autoscaler is supported for adjusting the replica count of deployments running on both virtual node pools and managed node pools.

B. The Kubernetes Cluster Autoscaler can be deployed to automatically resize both managed node pools and virtual node pools when pod resource demands change.

C. The Kubernetes Vertical Pod Autoscaler is restricted to managed node pools and is not supported for adjusting resources on virtual node pools.

D. The Kubernetes Metrics Server can be deployed as a standalone program to collect resource metrics on both virtual node pools and managed node pools.

Answer: (SHOW ANSWER)

Oracle explicitly lists Kubernetes Cluster Autoscaler among the deploymentsnot supported when using OKE virtual nodes. Virtual nodes provide a serverless data-plane model in which Oracle manages the underlying compute capacity, so customers do not resize a conventional virtual-node pool through Kubernetes Cluster Autoscaler. Cluster Autoscaler remains applicable to managed node pools, where worker-node capacity can be increased or decreased according to scheduling demand. Oracle also identifies Vertical Pod Autoscaler as unsupported on virtual nodes, while Metrics Server is supported through a deployment method adapted for virtual nodes. Horizontal Pod Autoscaler operates at the pod/deployment level and can adjust replica counts without requiring customers to resize virtual-node infrastructure. Therefore, B is the NOT valid statement because it incorrectly claims Cluster Autoscaler can resize both managed and virtual node pools.

NEW QUESTION: 14

In the DevOps lifecycle, what is the difference between continuous delivery and continuous deployment?

(Choose two.)

A. Continuous delivery involves automation of developer tasks, while continuous deployment involves manual operational tasks.
B. Continuous delivery utilizes automatic deployment to a development environment, while continuous deployment involves automatic deployment to a production environment.
C. Continuous delivery requires more automatic linting, while continuous deployment testing must be run manually.
D. Continuous delivery is a process that initiates deployment manually, while continuous deployment is based on automating the deployment process.
Answer: (SHOW ANSWER)
The two correct differences between continuous delivery and continuous deployment in the DevOps lifecycle are: Continuous delivery is a process that initiates deployment manually, while continuous deployment is based on automating the deployment process. In continuous delivery, the software is ready for deployment, but the decision to deploy is made manually by a human. On the other hand, continuous deployment automates the deployment process, and once the software passes all the necessary tests and quality checks, it is automatically deployed without human intervention. Continuous delivery involves automatic deployment to a development environment, while continuous deployment involves automatic deployment to a production environment. In continuous delivery, the software is automatically deployed to a development or staging environment for further testing and validation. However, the actual deployment to the production environment is performed manually. In continuous deployment, the software is automatically deployed to the production environment, eliminating the need for manual intervention in the deployment process. These differences highlight the level of automation and human involvement in the deployment process between continuous delivery and continuous deployment approaches in the DevOps lifecycle.

NEW QUESTION: 15
Which statement is NOT valid regarding custom metrics and the OCI Monitoring PostMetricData API?
A. The timestamp of any posted custom metric data point must be less than two hours in the past and less than 10 minutes in the future.
B. The custom metric namespace name specified in the payload must avoid using reserved system prefixes such as oci_ or oracle_.
C. The minimum supported ingestion frequency for posting raw custom metric data points to the service is one minute.
D. The PostMetricData API enforces a strict per-call limit of a maximum of 20 dimensions per metric group for the request.
Answer: (SHOW ANSWER)
Statement C is invalid because Oracle permits custom metrics to be postedas frequently as every second. The one-minute value applies to the minimum aggregation interval when metric data is subsequently queried, not to the ingestion frequency. Oracle also verifies the remaining statements. PostMetricData requires timestamps to be less than two hours in the past and less

than ten minutes in the future. Custom namespaces must not use reserved prefixes such as oci_ or oracle_. The API additionally allows a maximum of 20 dimensions per metric group in a single posting operation. Therefore, interpreting one minute as the minimum interval between raw metric submissions incorrectly confuses ingestion frequency with Monitoring ' s aggregation /query interval. Option C is consequently the NOT valid statement.

NEW QUESTION: 16
What can you use to dynamically make Kubernetes resources discoverable to public DNS servers? (Choose the best answer.)
A. kubeDNS
B. DynDNS
C. CoreDNS
D. ExternalDNS
Answer: (SHOW ANSWER)
To dynamically make Kubernetes resources discoverable to public DNS servers, you can use ExternalDNS.
ExternalDNS is a Kubernetes add-on that automates the management of DNS records for your Kubernetes services and ingresses. It can be configured to monitor the changes in your Kubernetes resources and automatically update DNS records in a supported DNS provider. By integrating ExternalDNS with your Kubernetes cluster, you can ensure that the DNS records for your services and ingresses are automatically created, updated, or deleted based on changes in your Kubernetes resources. This allows your Kubernetes resources to be discoverable by external systems through public DNS servers.

Valid 1z0-1084-26 Dumps shared by ExamDiscuss.com for Helping Passing 1z0-1084-26 Exam! ExamDiscuss.com now offer the **newest 1z0-1084-26 exam dumps**, the ExamDiscuss.com 1z0-1084-26 exam **questions have been updated** and **answers have been corrected** get the **newest** ExamDiscuss.com 1z0-1084-26 dumps with Test Engine here: https://www.examdiscuss.com/Oracle/exam/1z0-1084-26/premium/ (**182** Q&As Dumps, **35%OFF Special Discount Code: freecram**)

NEW QUESTION: 17
A DevOps engineer is troubleshooting the Meshifyd application, which is running in an Oracle Cloud Infrastructure (OCI) environment. The engineer has set up the OCI Logging service to store access logs for the application but notices that the logs from the Meshifyd application are not showing up in the logging service. The engineer suspects that there might be an issue with the logging configuration. Which two statements are potential reasons for logs from the Meshifyd application not showing up in the OCI Logging service?
A. The logconfig.json file has incorrect or missing OCID for the custom log in the logobjectId field.

B. The OCI Logging service is set up to pre access logs by creating a log group and custom log within the same compartment.
C. The logconfig.json file has incorrect or missing information in the application namespace in the paths field.
D. The logconfig.json file has incorrect or missing information in the application namespace in the src field.
E. The logconfig.json file has incorrect or missing OCID for the custom log group in the logGroupObjectId field.
Answer: (SHOW ANSWER)
The logconfig.json file is a configuration file that specifies how the Unified Monitoring Agent collects and uploads custom logs to the OCI Logging service2. The logconfig.json file contains an array of objects, each representing a custom log configuration2. Each custom log configuration object has the following fields2:
* logGroupObjectId: The OCID of the log group where the custom log is stored.
* logObjectId: The OCID of the custom log.
* paths: An array of paths to files or directories containing the custom logs.
* src: A regular expression that matches the files containing the custom logs.
* parser: A parser definition that specifies how to parse the custom logs. If the logconfig.json file has incorrect or missing OCID for the custom log in the logobjectId field, or incorrect or missing OCID for the custom log group in the logGroupObjectId field, then the Unified Monitoring Agent will not be able to upload the custom logs to the OCI Logging service2. Therefore, these are potential reasons for logs from the Meshifyd application not showing up in the OCI Logging service. Verified References: Unified Monitoring Agent Configuration File

NEW QUESTION: 18
Which query language is tailored specifically to request and search distributed tracing data in the OCI Application Performance Monitoring Trace Explorer?
A. Logging Query Language (LQL)
B. Monitoring Query Language (MQL)
C. Trace Query Language (TQL)
D. Analytics Query Language (AQL)
Answer: C (LEAVE A REPLY)
Oracle Application Performance Monitoring provides a dedicated query language in Trace Explorer for investigating distributed traces and their associated spans. Oracle documentation refers to it as theTrace Explorer Query Language, which uses a simplified SQL-like syntax with clauses such as show, filters, grouping, and aggregation. Among the supplied choices, this corresponds to Trace Query Language (TQL).
LQL belongs to OCI Logging and is used for searching log records, while MQL belongs to OCI Monitoring and operates on metric streams. AQL is not the Trace Explorer query language. Trace Explorer queries are specifically optimized for dimensions such as service names, operations, spans, traces, durations, and errors, making C the correct answer.

NEW QUESTION: 19
Your organization has deployed a custom OCI Authorizer Function in a compartment named secure-auth- compartment to perform token-based validation for all ingress traffic. You have set up an API Gateway in a compartment named public-gateway-compartment with an associated dynamic group named apigw-dg containing the gateway instance. The API Gateway deployment is currently returning HTTP 500 errors because the gateway lacks the necessary permissions to invoke the authorizer function during request validation. Which minimal IAM policy statement should you write to grant the API Gateway the required permission to successfully invoke the authorizer function?
A. Allow dynamic-group apigw-dg to use functions-family in compartment secure-auth-compartment
B. Allow dynamic-group apigw-dg to manage functions-family in compartment secure-auth-compartment
C. Allow dynamic-group apigw-dg to read functions-family in compartment public-gateway-compartment
D. Allow dynamic-group apigw-dg to inspect functions-family in compartment public-gateway-compartment
Answer: (SHOW ANSWER)
Oracle ' s documented API Gateway-to-Functions integration uses a dynamic group representing the API Gateway and an IAM policy granting that dynamic group use functions-family in the compartment containing the target Function. The use verb provides the operational permission necessary for invocation without granting unnecessary administrative control. manage functions-family would be excessively privileged because it includes lifecycle management capabilities. The read and inspect verbs are insufficient for executing a Function and, in options C and D, are also scoped to the wrong compartment. Because the authorizer resides in secure-auth-compartment, authorization must apply there. Consequently, A implements the Oracle-documented pattern and satisfies least privilege for API Gateway invocation of the custom authorizer Function.

NEW QUESTION: 20
You are developing a microservice that consumes messages from an Oracle Cloud Infrastructure (OCI) Queue and must delete each message after it is successfully processed. Which OCI CLI command should you use to delete an individual message using its receipt?
Freecram.net
A. oci queue messages delete --message-id < message_id > --queue-id < queue_OCID >
B. oci queue messages delete-message --message-id < message_id > --queue-id < queue_OCID >
C. oci queue messages delete --message-receipt < message_receipt > --queue-id < queue_OCID >
D. oci queue messages delete-message --message-receipt < message_receipt > --queue-id < queue_OCID >
Answer: (SHOW ANSWER)

OCI Queue requires the receipt returned when the message was obtained through GetMessages to delete that specific delivery. Oracle explicitly documents the command oci queue messages delete-message --message- receipt < message_receipt > --queue-id < queue_OCID > for deleting one message. A message ID alone is insufficient because deletion operates against the active message receipt associated with the current delivery and visibility state. Oracle provides the separate delete-messages operation for batch deletion and a queue purge operation for removing all messages. Therefore, options A and C use an invalid delete subcommand, while B incorrectly supplies --message-id. Option D exactly matches the current OCI CLI syntax and required receipt-based deletion mechanism.

NEW QUESTION: 21
Your organization mandates that all containerized workloads deployed to Oracle Cloud Infrastructure (OCI) Container Engine for Kubernetes (OKE) must utilize Kubernetes NetworkPolicy resources, such as Calico, to enforce strict micro-segmentation of traffic. You have deployed these workloads to an enhanced OKE cluster using virtual node pools to minimize infrastructure maintenance. During testing, you discover that the Calico installation fails because NetworkPolicy resources and third-party CNI policy integrations are completely unsupported on this specific node configuration. How should you resolve this security implementation challenge while continuing to use VCN-Native Pod Networking?
A. Reconfigure your virtual node pools to use the Flannel overlay plugin, which natively implements network policy isolation on virtual nodes.
B. Upgrade your OKE virtual nodes to Kubernetes version 1.30 or higher, which enables native, agentless NetworkPolicy support on virtual nodes.
C. Define an OCI IAM policy that grants the OKE workload identity access to manage network security lists at the virtual node level.
D. Migrate your workloads to managed node pools within the OKE cluster, where VCN-Native Pod Networking supports Calico for NetworkPolicy enforcement.
Answer: (SHOW ANSWER)
Oracle explicitly states that virtual nodes do not support Kubernetes NetworkPolicy resources or network providers such as Calico and Cilium that provide NetworkPolicy enforcement. Virtual nodes support only OCI VCN-Native Pod Networking and cannot use Flannel or other third-party CNI plugins. By contrast, Calico can be installed with managed node pools while retaining VCN-native pod networking. This combination enables Kubernetes NetworkPolicy objects to enforce pod-level traffic controls and micro- segmentation. An IAM policy or network security list cannot substitute for Kubernetes NetworkPolicy because OCI network controls and Kubernetes pod policy enforcement operate at different layers. Therefore, when Calico-based NetworkPolicy enforcement is mandatory, migrating the workload from virtual nodes to managed node pools is the supported OKE architecture.

NEW QUESTION: 22

A Docker image consists of one or more layers, each of which represents a Dockerfile instruction. The layers are stacked and each one is a delta of the changes from the previous layer. What permission is associated with these layers?

A. read mostly

B. write only

C. movable

D. read only

E. write once

Answer: (SHOW ANSWER)

The correct answer is: " read only. " The layers of a Docker image are read-only. Once a layer is created, it cannot be modified. Each layer represents a Dockerfile instruction, and it is stacked on top of the previous layer, forming a stack of immutable layers. These layers are designed to be read-only to ensure consistency and integrity of the image. When a Docker image is built, each instruction in the Dockerfile creates a new layer. Each layer represents the changes made by that instruction relative to the previous layer. The layers are stacked on top of each other to form the complete image. This layer-based approach allows for efficient storage and distribution of Docker images. Because the layers are read-only, any changes or modifications to the image result in the creation of new layers rather than modifying the existing ones. This immutability ensures that each layer remains intact and preserves the integrity of the image. It also enables Docker ' s caching mechanism, where previously built layers can be reused if the corresponding instructions haven ' t changed, speeding up the image build process. The other options mentioned, such as " write only, " " write once, " " movable, " and " read mostly, " do not accurately describe the permission associated with Docker image layers. Docker image layers are specifically designed to be read-only.

NEW QUESTION: 23

You are building a container image and pushing it to Oracle Cloud Infrastructure Registry (OCIR). You need to ensure that these images never get deleted from the repository. Which action should you take?

A. Create a group and assign a policy to perform lifecycle operations on images.

B. Write a policy to limit access to the specific repository in your compartment.

C. Edit the tenancy global retention policy.

D. Set global policy of image retention to " Retain All Images " .

Answer: (SHOW ANSWER)

To prevent images from being deleted from a repository, you can set the global policy of image retention to

"Retain All Images". This policy applies to all repositories in the tenancy and overrides any individual repository settings. You can also lock the policy to prevent any changes to it.Verified References: Using Retention Rules to Preserve Data

NEW QUESTION: 24

You are designing the persistence layer for a multi-service financial application where each microservice must completely isolate its private persistent data to enforce a bounded context. Which database configuration strategy should you implement within a multi-tenant Oracle Database to satisfy this requirement with the lowest operational overhead?

A. Assign a dedicated virtual machine DB system to each individual microservice.

B. Provision a single pluggable database (PDB) for each individual microservice.

C. Deploy a shared database schema with row-level security for all microservices.

D. Create a single container database (CDB) with shared tables for all microservices.

Answer: (SHOW ANSWER)

Oracle specifically recommends the multitenant database model for microservices because multiple Pluggable Databases can reside within one Container Database while retaining logical isolation. Oracle ' s microservices architecture guidance identifies PDBs as a strong persistence choice forbounded-context isolation, security, and high availability. Allocating a PDB per microservice isolates that service ' s private database objects while avoiding the infrastructure and operational overhead of provisioning an entire VM database system for every service. Shared schemas or shared tables weaken service autonomy and create tight coupling between bounded contexts. Because a CDB can host multiple isolated PDBs while sharing common database infrastructure and administration, option B best combines isolation with operational efficiency.

NEW QUESTION: 25

Which term describes a group formed by a master machine and a worker machine in a Kubernetes architecture?

A. Cluster

B. Node

C. Deployment

D. Container

E. Pod

Answer: (SHOW ANSWER)

The term that describes a group formed by a master machine and a worker machine in a Kubernetes architecture is " Cluster " . A cluster in Kubernetes consists of one or more master machines and multiple worker machines (also known as nodes). The master machine manages the overall control plane and orchestrates the deployment and management of containers on the worker nodes. The worker nodes are responsible for running the containers and executing the workloads. The cluster is the fundamental unit of organization and management in Kubernetes, providing the infrastructure and resources to run and manage containerized applications. It ensures high availability, scalability, and fault tolerance for the applications deployed within it.

NEW QUESTION: 26

Which is ONE of the differences between a microservice and a serverless function?

A. Microservices are used for long running operations while serverless functions are used for short running operations.

B. Microservices are triggered by events while serverless functions are not.
C. Microservices are stateless while serverless functions are stateful.
D. Microservices always use a data store while serverless functions never use a data store.
Answer: (SHOW ANSWER)
The correct answer is: Microservices are used for long running operations while serverless functions are used for short running operations. One of the key differences between microservices and serverless functions is the duration of their execution. Microservices are typically designed to handle long-running operations and may continuously run and process requests as part of a larger system. They are often deployed and managed as long-lived services. On the other hand, serverless functions are designed to handle short-lived operations or tasks that execute in response to specific events or triggers. They are event-driven and execute only when invoked, providing a lightweight and ephemeral computing model. Serverless functions are often used for executing small, isolated pieces of code without the need for managing infrastructure or scaling concerns.
While both microservices and serverless functions can be stateless or stateful depending on the specific implementation, the key distinction lies in the typical duration and execution pattern of these components within an application architecture.

NEW QUESTION: 27
You are setting up OCI Application Performance Monitoring (APM) to trace transactions in a Spring Boot application. Which installer command should you execute to provision the APM Java Agent with its target application service name, destination directory, and credentials obtained from your APM domain?
A. java -jar apm-java-agent-installer.jar register-agent -destination=/opt/apm -app-name=springboot-app - private-data-key=KeyA123 -data-upload-endpoint=https://endpoint.com
B. java -jar apm-java-agent-installer.jar install-agent -path=/opt/apm -service-name=springboot-app - upload-key=KeyA123 -upload-endpoint=https://endpoint.com
C. java -jar apm-java-agent-installer.jar configure-agent -destination=/opt/apm -service-name=springboot- app -public-data-key=KeyA123 -data-upload-endpoint=https://endpoint.com
D. java -jar apm-java-agent-installer.jar provision-agent -destination=/opt/apm -service-name=springboot- app -private-data-key=KeyA123 -data-upload-endpoint=https://endpoint.com
Answer: (SHOW ANSWER)
Oracle APM documentation specifies the provision-agent operation for installing and configuring the APM Java Agent. Its mandatory arguments include -service-name, identifying the monitored service; -destination, specifying where the agent is provisioned; -private-data-key, using the installation key generated for the APM domain; and -data-upload-endpoint, identifying the domain ' s upload endpoint. Oracle provides essentially the same command structure as option D in its provisioning documentation. register-agent, install-agent, and configure-agent are not the documented provisioning subcommands, while the alternative parameter names such as -app-name, -upload-key, and -path do not match Oracle ' s mandatory syntax. Consequently, option D is the fully documented command.

NEW QUESTION: 28
Which TWO are part of the Cloud Native Computing Foundation (CNCF) container runtime? (Choose two.)
A. rkt-o
B. runc
C. getcd
D. containerd
Answer: (SHOW ANSWER)
The two components that are part of the Cloud Native Computing Foundation (CNCF) container runtime are:
containerd: containerd is an open-source container runtime that provides a runtime environment for containers, including managing container images, executing containers, and handling container lifecycle events. It is designed to be lightweight and extensible, providing the necessary functionality to run containers efficiently. runc: runc is a lightweight container runtime that serves as a reference implementation of the Open Container Initiative (OCI) runtime specification. It is responsible for launching and managing containers based on OCI specifications, including handling container isolation, namespaces, cgroups, and other low-level container operations. These two components, containerd and runc, are widely used in the container ecosystem and are part of the CNCF ' s efforts to promote and develop open-source technologies for cloud-native computing.

NEW QUESTION: 29
Which statement is NOT valid regarding the comparison between cloud-native and cloud-based applications?
A. Cloud-native applications are built specifically for the cloud to minimize dependencies, whereas cloud- based applications are legacy systems modified to run in a cloud environment.
B. Cloud-native applications utilize a microservices architecture to reduce the risk of downtime, whereas cloud-based applications are often more vulnerable to central failure points.
C. Cloud-native applications require manual host provisioning to scale their compute resources, whereas cloud-based applications scale automatically and dynamically by default.
D. Cloud-native applications leverage pay-as-you-go pricing to optimize resource costs, whereas cloud- based applications can incur unnecessary expenses due to over-provisioning.
Answer: (SHOW ANSWER)
Option C reverses a fundamental cloud-native characteristic. Oracle describes cloud-native applications as being designed for automated scale-up and scale-down in response to changing workload demand.
Technologies including containers, Kubernetes, serverless architectures, microservices, and automated DevOps processes are specifically intended to reduce manual infrastructure management. Oracle also identifies scalability, resilience, reduced operational overhead, and cost optimization through cloud consumption models as major cloud-native advantages. Applications

migrated to cloud infrastructure without substantial architectural modernization can retain monolithic dependencies, fixed capacity assumptions, and manual provisioning practices. Therefore, claiming that cloud-native applications require manual host provisioning while cloud-based applications automatically scale by default is technically inverted. C is consequently the NOT valid statement among the available choices.

NEW QUESTION: 30

You have two microservices, A and B, running in production. Service A relies on APIs from service B. You want to test changes to service A without deploying all of its dependencies, which include service B. Which approach should you take to test service A?

Freecram.net

A. Test using API mocks.
B. Test the APIs in private environments.
C. Test against production APIs.
D. There is no need to explicitly test APIs.
Answer: (SHOW ANSWER)
API mocking is a technique that simulates the behavior of real APIs without requiring the actual implementation or deployment of the dependent services1. API mocking allows you to test changes to service A without deploying all of its dependencies, such as service B, by creating mock responses for the APIs that service A relies on1. API mocking has several benefits, such as1:
* Faster testing: You can test your service A without waiting for service B to be ready or available, which reduces the testing time and feedback loop.
* Isolated testing: You can test your service A in isolation from service B, which eliminates the possibility of external factors affecting the test results or causing errors.
* Controlled testing: You can test your service A with different scenarios and edge cases by creating mock responses that mimic various situations, such as success, failure, timeout, etc.

NEW QUESTION: 31
Which statement about microservices is FALSE?
A. They are typically designed around business capabilities.
B. It is fairly common for them to communicate with HTTP.
C. Multiple microservices can run in one process.
D. They are independently deployable.
Answer: (SHOW ANSWER)
The correct answer is: " Multiple microservices can run in one process. " The statement that is FALSE about microservices is: " Multiple microservices can run in one process. " Microservices are designed to be independent and autonomous services that can be deployed, scaled, and managed separately. Each microservice typically runs as its own separate process. This separation allows for loose coupling and independent scalability of individual services. Running multiple microservices in a single process would violate the fundamental principle of microservices architecture, which emphasizes the isolation and independence of services.

Combining multiple microservices into a single process would tightly couple them, making it difficult to independently manage, scale, and update each service. The other statements provided are true about microservices: It is fairly common for microservices to communicate with HTTP: Microservices often communicate with each other using lightweight protocols like HTTP/REST. This allows for interoperability, flexibility, and ease of communication between microservices. Microservices are typically designed around business capabilities: Microservices architecture advocates designing services based on specific business capabilities or functionalities. Each microservice focuses on a specific business domain or task, enabling modular and scalable development. They are independently deployable: Microservices are designed to be independently deployable units. This means that each microservice can be developed, tested, deployed, and updated separately without affecting other services. This flexibility allows for faster development cycles and more efficient scalability.

Valid 1z0-1084-26 Dumps shared by ExamDiscuss.com for Helping Passing 1z0-1084-26 Exam! ExamDiscuss.com now offer the **newest 1z0-1084-26 exam dumps**, the ExamDiscuss.com 1z0-1084-26 exam **questions have been updated** and **answers have been corrected** get the **newest** ExamDiscuss.com 1z0-1084-26 dumps with Test Engine here: https://www.examdiscuss.com/Oracle/exam/1z0-1084-26/premium/ (**182** Q&As Dumps, **35%OFF Special Discount Code: freecram**)

NEW QUESTION: 32
You are developing an asynchronous, decoupled billing microservice that consumes payment messages from an Oracle Cloud Infrastructure (OCI) Queue. You notice that some corrupted payment payloads are causing the consumer application to fail and restart repeatedly, leaving the messages unprocessed and blocking the queue. To isolate these toxic messages, you want to redirect them to a Dead Letter Queue (DLQ) after 5 failed delivery attempts.
Which TWO configuration steps or characteristics must you apply to implement this solution?
A. Modify the " Maximum delivery attempts " setting of the queue to 5.
B. A separate target queue must be manually created first and then selected as the designated Dead Letter Queue.
C. A dead letter queue is automatically created when you create the queue, so no manual DLQ resource creation is required.
D. Configure an IAM policy specifically allowing the source queue to write to a target DLQ queue.
E. Enable consumer groups on the queue because DLQ functionality is only supported when utilizing consumer groups.
Answer: (SHOW ANSWER)
The supplied Word answer key requires correction here. Current Oracle Queue documentation explicitly states that adead letter queue is automatically created when a queue is created. Therefore, option C is correct and option B is not. Oracle also provides the Maximum delivery

attempts setting, which controls how many unsuccessful deliveries can occur before a problematic message is transferred to the DLQ. Setting that value to 5 implements the requested retry threshold. No separately provisioned target queue is required, consumer groups are optional and unrelated to enabling DLQ functionality, and administrators do not create an IAM policy authorizing a source queue resource to write to its automatically associated DLQ. Consequently, the verified current answer isA, C, not the Word file ' s A, B.

NEW QUESTION: 33
Which TWO statements accurately describe an Oracle Functions application? (Choose two.)
A. A common context to store configuration variables that are available to all functions in the application.
A Docker image containing all the functions that share the same configuration.
B. An application based on Oracle Functions, Oracle Cloud Infrastructure (OCI) Events, and OCI API Gateway services.
C. A small block of code invoked in response to an OCI Events service.A logical group of functions.
D. A Docker image containing all the functions that share the same configuration.
Answer: (SHOW ANSWER)
The correct statements are: A common context to store configuration variables that are available to all functions in the application. A Docker image containing all the functions that share the same configuration. A logical group of functions.Explanation: An Oracle Functions application provides a common context for functions within the application. It allows you to store configuration variables that are accessible by all the functions in the application. Functions within the same application can share the same Docker image, which contains the common configuration and dependencies. An Oracle Functions application serves as a logical group that organizes related functions. Functions within the same application can be managed collectively, and they can interact and share resources within the application context.

NEW QUESTION: 34
When creating an alarm query in the OCI Monitoring service, which MQL trigger rule evaluates the lack of incoming raw metric data points as the trigger criteria?
A. absence
B. threshold
C. missing
D. non-zero
Answer: (SHOW ANSWER)
OCI Monitoring provides anabsencetrigger for alarms that must fire when a metric stops emitting data. Oracle
' s current documentation states that an absence alarm uses the MQL absent() statistic. It returns 1 when the metric is absent for the complete evaluation interval and 0 when data is present. An expression can therefore take the form CpuUtilization[1m].absent() or use groupBy() before

absent() when separate dimensional streams must be evaluated. A threshold trigger, by contrast, evaluates numeric metric values against limits such as greater than 90 percent. missing and non-zero are not OCI Monitoring trigger-rule names for the required behavior. Thus, the correct trigger rule for detecting a lack of incoming metric data is absence.

NEW QUESTION: 35
You have two microservices, A and B running in production. Service A relies on APIs from service B. You want to test changes to service A without deploying all of its dependencies, which includes service B. Which approach should you take to test service A?
A. Test using a previous test version of service B.
B. Test using an API mock of service B.
C. Test using the current production version of service B.
D. This is not possible because service B is a dependency.
Answer: (SHOW ANSWER)
The correct answer is: Test using an API mock of service B. To test service A without deploying all of its dependencies, including service B, you can use an API mock of service B. An API mock is a simulated version of the API that mimics the behavior of the actual service. By using an API mock, you can isolate the testing of service A and simulate the responses and behavior of service B ' s APIs. With an API mock, you can define the expected responses and behavior of service B ' s APIs, allowing you to test the integration between service A and the mocked version of service B. This approach enables you to verify the functionality of service A without relying on the availability or changes in the actual service B. By decoupling the dependencies and using an API mock, you can perform independent testing of service A, ensuring its functionality in isolation.

NEW QUESTION: 36
Which limit represents the maximum write throughput supported per partition of an Oracle Cloud Infrastructure stream when designing high-volume data producer workloads?
A. 1 MB per second
B. 2 MB per second
C. 4 MB per second
D. 8 MB per second
Answer: (SHOW ANSWER)
Oracle ' s current OCI Streaming limits specify that each stream partition supports a total data write rate of 1 MB per second. Partition count is therefore a fundamental capacity-planning parameter for producer workloads. If an application is expected to produce more than 1 MB per second, the stream must contain enough partitions to distribute the required ingress throughput appropriately. Oracle ' s partitioning guidance likewise states that expected throughput should be calculated before stream creation because each partition is limited to the 1 MB-per-second write rate. The 2 MB figure is not the documented write limit, while 4 MB and 8 MB exceed the per-partition ingress capacity. Therefore, the maximum write throughput represented by the available choices is 1 MB per second.

NEW QUESTION: 37
You are instructed to automate manual tasks and help software teams manage complex environments at scale using the Oracle Cloud Infrastructure (OCI) services. Which THREE OCI services can be leveraged to securely store and version your application ' s source code, and automate the building, testing, and deployment of applications to the OCI platform? (Choose three.)
A. DevOps
B. Container Engine for Kubernetes
C. Oracle APEX Application Development
D. Resource Manager
E. Oracle Cloud Infrastructure Registry
F. Oracle Cloud Logging Analytics
Answer: (SHOW ANSWER)
The three OCI services that can be leveraged to securely store and version your application ' s source code, and automate the building, testing, and deployment of applications to the OCI platform are: DevOps: OCI provides a comprehensive set of DevOps services, including Oracle Developer Cloud Service, which allows you to manage source code repositories, automate builds and testing, and streamline the deployment process.
Container Engine for Kubernetes: OCI ' s Container Engine for Kubernetes (OKE) enables you to deploy and manage containerized applications using Kubernetes. It provides a scalable and reliable platform for automating the deployment of your applications. Oracle Cloud Infrastructure Registry: OCI Registry is a fully managed, private container registry that allows you to securely store and manage Docker images. It integrates with other OCI services, such as Container Engine for Kubernetes, to facilitate seamless deployment and orchestration of containerized applications. These services combined provide the necessary tools and infrastructure to support continuous integration and continuous deployment (CI/CD) workflows, enabling efficient and automated application development and deployment processes in the Oracle Cloud Infrastructure environment.

NEW QUESTION: 38
Your cloud-native banking microservice retrieves API tokens stored in OCI Vault. To meet strict security compliance standards, you must configure secret rules on these vault secrets. The security policies mandate that the secret contents must automatically expire and have their retrieval blocked 30 days after creation, and any previously used token content must never be allowed for reuse, even if the older secret version containing that content has been deleted.
Which TWO configurations must you specify within the secret rules to satisfy these compliance requirements?
A. Set up a SECRET_EXPIRY_RULE with secret_version_expiry_interval set to " P30D " and is_secret_content_retrieval_blocked_on_expiry set to false.

B. Set up an HSM_PURGE_RULE that automatically runs a background execution task to permanently erase deleted secrets from the vault hardware.
C. Set up a SECRET_EXPIRY_RULE with secret_version_expiry_interval set to " P30D " and is_secret_content_retrieval_blocked_on_expiry set to true.
D. Set up a SECRET_REUSE_RULE with is_enforced_on_deleted_secret_versions set to true.
E. Set up a SECRET_REUSE_RULE with is_enforced_on_deleted_secret_versions set to false.
Answer: (SHOW ANSWER)
OCI Vault provides both Secret Expiry Rules and Secret Reuse Rules.
secret_version_expiry_interval uses an ISO 8601 duration, so P30D establishes a 30-day validity period. Setting is_secret_content_retrieval_blocked_on_expiry to true ensures secret contents cannot be retrieved after expiry, satisfying the first requirement. For the second requirement, a SECRET_REUSE_RULE prevents previously used secret contents from being reused across versions. Oracle ' s SDK documentation confirms that is_enforced_on_deleted_secret_versions=true extends that restriction even when the previous secret version has been deleted. Setting either Boolean to false would weaken the stated controls.
HSM_PURGE_RULE is not one of the documented OCI Vault secret-rule types. Therefore, C and D jointly satisfy both requirements.

NEW QUESTION: 39
What are the TWO main reasons you would choose to implement a serverless architecture? (Choose two.)
A. No need for integration testing
B. Automatic horizontal scaling
C. Easier to run long-running operations
D. Reduced operational cost
E. Improved in-function state management
Answer: (SHOW ANSWER)
The two main reasons to choose a serverless architecture are: Automatic horizontal scaling: Serverless architectures allow for automatic scaling of resources based on demand. The infrastructure automatically provisions and scales resources as needed, ensuring that applications can handle varying workloads efficiently. This eliminates the need for manual scaling and optimizes resource utilization. Reduced operational cost: Serverless architectures follow a pay-per-use model, where you are billed only for the actual execution time and resources consumed by your functions. This leads to cost savings as you don ' t have to pay for idle resources. Additionally, serverless architectures remove the need for managing and maintaining servers, reducing operational overhead and associated costs.Explanation: No need for integration testing:
Integration testing is still necessary in serverless architectures to ensure that functions integrate correctly with other components and services. Serverless functions can interact with various

event sources, databases, and APIs, and testing is required to verify the integration points.
Improved in-function state management:
Serverless architectures typically encourage stateless functions that operate on short-lived requests or events.
While there are mechanisms to manage state within a function, serverless architectures are designed to be stateless by default, promoting scalability and fault tolerance. Easier to run long-running operations:
Serverless functions are generally designed for short-lived operations rather than long-running tasks. If you have a requirement for long-running operations, a serverless architecture may not be the ideal choice, as it has execution time limits and may not provide the necessary resources for extended execution.

NEW QUESTION: 40

Which technique is used for testing the entire user flow as well as the moving parts of a cloud native app, ensuring that there are no high-level discrepancies?

Freecram.net

A. Contract Testing
B. Integration Testing
C. Unit Testing
D. Component Testing
E. End-to-end Testing
Answer: (SHOW ANSWER)
End-to-end testing is a technique that involves checking the entire user flow as well as the moving parts of a cloud native app, ensuring that there are no high-level discrepancies3. End-to-end testing simulates real user scenarios and validates the functionality, performance, reliability, and security of the app from start to finish3. End-to-end testing has several benefits, such as3:
* Comprehensive testing: You can test your app as a whole and verify that all the components work together as expected.
* User-centric testing: You can test your app from the user's perspective and ensure that it meets the user' s needs and expectations.
* Quality assurance: You can test your app in a realistic environment and identify any issues or defects before releasing it to the users.

NEW QUESTION: 41
What is the maximum execution time of Oracle Functions?
A. 240 seconds
B. 300 seconds
C. 60 seconds
D. 120 seconds
Answer: (SHOW ANSWER)
The maximum execution time of Oracle Functions is 300 seconds, which is equivalent to 5 minutes. This means that a function running within Oracle Functions cannot exceed a runtime of 5

minutes. If a function requires longer execution times, alternative approaches such as invoking external services asynchronously or using long-running processes should be considered. It is important to design functions with this execution time limitation in mind to ensure optimal performance and efficiency within the Oracle Functions platform.

NEW QUESTION: 42
You have created a repository in Oracle Cloud Infrastructure Registry (OCIR) in the us-ashburn-1 (iad) region in your tenancy with the heyoci namespace. Which three are valid tags for an image named myapp? (Choose three.)
A. iad.ccir.io/heyoci/myapp:latest
B. iad.ocir.io/heyoci/myapp:0.0.2-beta
C. iad.ocir.io/heyoci/myproject/myapp:0.0.1
D. us-ashburn-1.ocir.io/myproject/heyoci/myapp:latest
E. us-ashburn-1.ocir.io/heyoci/myapp:0.0.2-beta
F. us-ashburn-1.ocir.io/heyoci/myproject/myapp:0.0.2-beta
Answer: (SHOW ANSWER)
(OPT) iad.ocir.io/myproject/heyoci/myapp:latest The three valid tags for an image named myapp in the heyoci namespace of the Oracle Cloud Infrastructure Registry (OCIR) in the us-ashburn-1 (iad) region are:
iad.ccir.io/heyoci/myapp:latest This tag specifies the latest version of the myapp image in the heyoci namespace in the iad region. iad.ocir.io/heyoci/myapp:0.0.2-beta This tag specifies version 0.0.2-beta of the myapp image in the heyoci namespace in the iad region. us-ashburn-1.ocir.io/heyoci/myapp:0.0.2-beta This tag specifies version 0.0.2-beta of the myapp image in the heyoci namespace in the us-ashburn-1 region. The other options mentioned are not valid tags for the given scenario: us-ashburn-1.ocir.io/heyoci/myapp:0.0.2- beta (incorrect format, missing the namespace) iad.ocir.io/myproject/heyoci/myapp:latest (incorrect format, incorrect order of namespace and project) us-ashburn-1.ocir.io/myproject/heyoci/myapp:latest (incorrect format, incorrect order of namespace and project) The correct format for tagging an image in OCIR is < region > .ocir.io/ < namespace > / < repository > : < tag > .

NEW QUESTION: 43
You are deploying a microservices-based application to an Oracle Cloud Infrastructure Container Engine for Kubernetes cluster running a modern Kubernetes release. You need to identify the underlying container runtime interface implementation that manages and runs your Open Container Initiative-compliant images within the worker nodes.
Which container runtime interface is utilized by the cluster for this purpose?
A. containerd
B. Docker Engine
C. rkt runtime
D. CRI-O
Answer: (SHOW ANSWER)

OKE transitioned from Docker toCRI-Obeginning with Kubernetes 1.20.8. Oracle documents CRI-O as its Kubernetes Container Runtime Interface implementation for modern OKE worker nodes. CRI-O is purpose- built for Kubernetes and runs Open Container Initiative-compliant images, so existing Docker-format images remain usable even though Docker Engine itself is no longer the underlying Kubernetes runtime. Oracle ' s current OKE documentation continues to reference CRI-O on modern node pools, including current releases where runtime configuration changes can affect container behavior. containerd is common in other Kubernetes distributions but is not the runtime identified by Oracle for these OKE nodes. rkt is obsolete, and Docker Engine was used only by older OKE versions. Therefore, D is correct.

NEW QUESTION: 44
Which of these is NOT a valid authentication method for accessing an OCI API Gateway deployment?
A. HTTP Basic
B. API Key
C. OAuth
D. SAML Token
Answer: (SHOW ANSWER)
OCI API Gateway supports the following authentication methods for accessing an API deployment3:
* HTTP Basic: The client sends a username and password with each request. The credentials are validated against a user database in Oracle Identity Cloud Service (IDCS).
* API Key: The client sends an API key with each request. The API key is validated against a list of keys stored in IDCS or OCI Vault.
* OAuth: The client obtains an access token from an authorization server (such as IDCS) and sends it with each request. The access token is validated against the authorization server and optionally checked for required scopes.
* JWT Token: The client obtains a JSON Web Token (JWT) from an identity provider (such as IDCS or OCI IAM) and sends it with each request. The JWT is validated against the identity provider's public key and optionally checked for required claims. SAML Token is not a valid authentication method for accessing an OCI API Gateway deployment. SAML is an XML-based standard for exchanging authentication and authorization data between different parties, such as a service provider and an identity provider4. SAML tokens are typically used for web browser single sign-on (SSO) scenarios, not for API access4.

NEW QUESTION: 45
You are configuring an Oracle Cloud Infrastructure Events rule to process lifecycle changes for both Object Storage objects and Compute instances within a single compartment. Which JSON filter attribute structure correctly groups multiple event types under a single filter condition?
A. " actionTypes " : " com.oraclecloud.objectstorage.createobject | com.oraclecloud.computeapi.instance.

terminate.begin "
B. " eventTypes " : " com.oraclecloud.objectstorage.createobject,
com.oraclecloud.computeapi.instance.
terminate.begin "
C. " filterType " : [" com.oraclecloud.objectstorage.createobject " , " com.oraclecloud.computeapi.
instance.terminate.begin "]
D. " eventType " : [" com.oraclecloud.objectstorage.createobject " , "
com.oraclecloud.computeapi.
instance.terminate.begin "]
Answer: (SHOW ANSWER)
Oracle Events uses the singular top-level field eventType to identify the type of event represented by an event message. When a rule must match more than one event type, Oracle explicitly documents placing multiple event-type strings in a JSON array assigned to eventType. Values within that array operate as alternatives, allowing the condition to match an event corresponding to any listed type. eventTypes, actionTypes, and filterType are not valid substitutes for the documented field. Likewise, concatenating event types into a single pipe-delimited or comma-delimited string does not create multiple event matches. Option D therefore implements the exact structure Oracle shows in its Events filtering examples for a rule covering multiple event types.

NEW QUESTION: 46
What is the difference between blue/green and canary deployment strategies? (Choose the best answer.)
A. In blue/green, current applications are slowly replaced with new ones. In canary, the application Is deployed Incrementally to a select group of people.
B. In blue/green, both old and new applications are in production at the same time. In canary, the application Is deployed incrementally to a select group of people.
C. In blue/green, current applications are slowly replaced with new ones. In canary, both old and new applications are in production at the same time.
D. In blue/green, the application Is deployed In minor Increments to a select group of people. In canary, both old and new applications are simultaneously in production.
Answer: (SHOW ANSWER)
The correct answer is: In blue/green deployment, both old and new applications are in production at the same time. In canary deployment, the application is deployed incrementally to a select group of people. In a blue
/green deployment strategy, two identical environments, referred to as blue and green, are set up. The current production environment (blue) continues to serve live traffic while a new version of the application is deployed in the green environment. Once the new version is tested and deemed stable, traffic is routed from the blue environment to the green environment, making it the new production environment. This approach allows for a seamless switch between the old and new versions of the application. On the other hand, in a canary deployment strategy, the new version of the application is deployed incrementally to a small subset of users or a specific group. This

allows for testing the new version in a real production environment while minimizing the impact of any potential issues. If the new version performs well and meets the desired criteria, it can be gradually rolled out to a larger audience or the entire user base. In summary, the main difference between blue/green and canary deployment strategies lies in how the deployment is managed. Blue/green involves simultaneous production of both old and new applications, while canary deployment focuses on incremental deployment to a select group of users.

Valid 1z0-1084-26 Dumps shared by ExamDiscuss.com for Helping Passing 1z0-1084-26 Exam! ExamDiscuss.com now offer the **newest 1z0-1084-26 exam dumps**, the ExamDiscuss.com 1z0-1084-26 exam **questions have been updated** and **answers have been corrected** get the **newest** ExamDiscuss.com 1z0-1084-26 dumps with Test Engine here: https://www.examdiscuss.com/Oracle/exam/1z0-1084-26/premium/ (**182** Q&As Dumps, **35%OFF Special Discount Code: freecram**)

NEW QUESTION: 47
Your Oracle Cloud Infrastructure (OCI) Container Engine for Kubernetes (OKE) administrator has created an OKE cluster with one node pool in a public subnet. You have been asked to provide a log file from one of the nodes for troubleshooting purpose. Which step should you take to obtain the log file?
A. Use the username opc and password to login.
B. It is impossible because OKE is a managed Kubernetes service.
C. SSH into the nodes using the private key.
D. SSH into the node using the public key.
Answer: (SHOW ANSWER)
To obtain a log file from one of the nodes in an Oracle Cloud Infrastructure (OCI) Container Engine for Kubernetes (OKE) cluster, you should SSH into the nodes using the private key. Here ' s the step-by-step process: Obtain the private key: The private key is required to authenticate and access the nodes in the OKE cluster. You should obtain the private key from your administrator or the appropriate key pair used to create the cluster. SSH into the node: Use a secure shell (SSH) client, such as OpenSSH, to connect to the desired node in the cluster. The SSH command typically includes the private key file path and the public IP address or hostname of the node. Example command: ssh -i < private_key_file > opc@ < node_public_ip > Replace < private_key_file > with the path to the private key file and < node_public_ip > with the public IP address of the node you want to access. Navigate to the log file location: Once you have successfully connected to the node, navigate to the directory where the log file is located. The exact location and name of the log file may vary depending on the Kubernetes distribution and configuration. Copy or view the log file: You can either copy the log file from the node to your local machine using the scp command or view the contents directly on the node using tools like cat or

less. By following these steps, you will be able to access the log file from the desired node in the OKE cluster for troubleshooting purposes.

NEW QUESTION: 48
Your development team wants to configure API Gateway request validation to shield an upstream backend microservice from ill-formed API calls. You want to test the validation rules in your staging environment first without blocking any non-compliant client requests, and you want to ensure the gateway checks that a mandatory HTTP header is present. Which TWO configurations should you apply to your API Gateway deployment specification to meet this requirement?
A. Set the request validation policy mode to Permissive to log validation failures without rejecting requests.
B. Set the request validation policy mode to Enforcing to forward failed requests to the execution log instead of blocking them.
C. Configure the request validation policy with header validations to mandate the presence of the required HTTP header.
D. Create a dynamic group and IAM policy to grant the API Gateway permission to inspect incoming request bodies in real-time.
E. Use the request.auth[< key >] context variable inside a Stock Response backend to log non-compliant headers.
Answer: (SHOW ANSWER)
OCI API Gateway request validation supportsEnforcing, Permissive, and Disabledvalidation modes. In Permissive mode, requests are validated and failures are recorded, but even non-compliant requests continue to the backend, making it appropriate for testing the impact of validation before enforcing it. Oracle also supports explicit header validation rules in which a header can be marked required: true. Enforcing mode instead blocks requests that fail validation and returns a 4xx response. No dynamic group is required for API Gateway to validate standard HTTP headers, and request.auth relates to authentication context rather than header-presence validation. Therefore, A and C directly satisfy both stated requirements.

NEW QUESTION: 49
Your enterprise is migrating a synchronous order fulfillment system to a distributed microservices architecture on Oracle Cloud Infrastructure. The development team originally implemented a Two-Phase Commit protocol to handle transactions spanning multiple databases, but it caused severe performance bottlenecks due to extended resource locking periods. You must now design a decentralized transaction strategy that processes steps as a sequence of independent local transactions, publishes events to coordinate subsequent steps, and executes compensating actions if a step fails. Which transaction design pattern should you implement to meet these requirements?
A. Two-Phase Commit Pattern
B. Transactional Outbox Pattern
C. Saga Design Pattern

D. Event Aggregation Pattern

Answer: (SHOW ANSWER)

The Saga Design Pattern is specifically intended for distributed transactions across microservices without the extended locking associated with Two-Phase Commit. Oracle explains that a Saga decomposes a business transaction into a sequence of independent local transactions. After completing its local operation, each service publishes an event or message that triggers the next participant. If a subsequent step fails, corresponding compensating transactions reverse previously completed operations and restore consistent business state. This approach shortens lock duration and supports eventual consistency across separate databases. Transactional Outbox solves reliable database-update-and-message-publishing coordination but does not by itself manage the entire distributed transaction sequence. Therefore, the required pattern is the Saga Design Pattern.

NEW QUESTION: 50

You are planning the migration of a legacy monolithic billing application to a cloud-native architecture on Oracle Cloud Infrastructure. The application currently relies on a centralized relational database and synchronous local procedure calls, and you must address the primary architectural challenges associated with transitioning to a distributed microservices model.
Which TWO statements represent valid engineering challenges or design strategies that must be considered during this cloud-native migration?

Freecram.net

A. Implementing microservices requires adopting eventual consistency patterns, such as the Saga pattern, because distributed databases cannot easily maintain synchronous transactions across service boundaries.

B. Transitioning to microservices increases the operational complexity of the application, necessitating centralized logging and distributed tracing to troubleshoot service-to-service communication.

C. By default, OCI ensures that all distributed containerized microservices automatically share a single, global transaction state across different compartments without any application-level coding.

D. Monolithic databases migrated to OCI Object Storage will automatically scale horizontally and support distributed transactions natively for all microservices.

E. Cloud-native microservices must communicate using synchronous REST APIs exclusively, as asynchronous event-driven communication is not supported within OCI.

Answer: (SHOW ANSWER)

Oracle identifies distributed transaction consistency and observability as major concerns when moving from monoliths to microservices. The Saga pattern addresses cross-service business transactions by dividing them into local transactions and using events plus compensating actions, avoiding long-duration distributed locks.

Oracle also emphasizes unified observability for microservices, including metrics, logs, and traces, because failures become harder to diagnose when calls cross multiple networked services. OCI does not automatically create a global application transaction across compartments, Object Storage is not a transactional relational database replacement, and

microservices are not restricted to synchronous REST communication. Oracle explicitly supports asynchronous messaging and event-driven designs for loose coupling. Therefore, A and B correctly describe realistic migration challenges and strategies.

NEW QUESTION: 51
You are developing a diagnostic tool to programmatically extract error events from multiple application custom logs in a specific compartment. Which OCI CLI command should your tool run to submit this log search query for a designated time window?
A. oci logging log search-logs --search-query < query > --time-start < start_time > --time-end < end_time >
B. oci logging-search search-logs --search-query < query > --time-start < start_time > --time-end < end_time >
C. oci logging search-logs --query-string < query > --time-start < start_time > --time-end < end_time >
D. oci logging-query search-logs --search-query < query > --time-start < start_time > --time-end < end_time >
Answer: B (LEAVE A REPLY)
Oracle provides a dedicated logging-search OCI CLI command group for querying indexed log content. The documented syntax is oci logging-search search-logs, with --search-query, --time-start, and --time-end as the principal parameters required to define the Logging Query Language expression and search interval. The ordinary oci logging command group is used for administration of logs and log groups rather than searching their contents. logging-query is not a valid top-level OCI CLI service group, and --query-string is not the documented search parameter. Therefore, option B exactly matches Oracle ' s current CLI command reference for submitting a programmatic log search across custom or service logs.

NEW QUESTION: 52
You are deploying a set of Java microservices to an Oracle Cloud Infrastructure Container Engine for Kubernetes cluster. You must ensure that the application design strictly adheres to the Twelve-Factor App methodology for configuration management and backing services to simplify multi-environment deployments. Which TWO design strategies must you implement to satisfy these Twelve-Factor App principles?
A. Inject all environment-specific configuration values directly into the compiled application code bundle during the build phase.
B. Externalize configuration values by utilizing environment variables or Kubernetes ConfigMaps to override default microservice settings.
C. Package each database credential and external API endpoint URL inside the application ' s local properties file for strict encapsulation.
D. Treat backing databases and message brokers as attached resources, allowing them to be swapped via configuration changes without modifying microservice code.

E. Maintain a single database schema instance that is shared dynamically across all environments to ensure absolute dev/prod database state parity.

Answer: (SHOW ANSWER)

The Twelve-Factor methodology requires configuration to be separated from application code and backing services to be treated as attached resources. Oracle ' s Twelve-Factor guidance specifically recommends keeping configuration outside the Docker image and supplying deployment-specific values through the environment. This allows the same application artifact to move between development, test, and production without rebuilding it. Kubernetes ConfigMaps or environment variables implement this principle effectively.

The backing-services factor similarly treats databases, queues, caches, and external services as replaceable resources referenced through configuration rather than hardcoded dependencies. Consequently, switching a database or message broker should require configuration changes rather than application-code changes.

Hardcoding credentials or environment-specific values violates configuration separation, while sharing one database across environments is not required for dev/prod parity. Therefore, B and D are correct.

Top of Form

Bottom of Form

NEW QUESTION: 53

Your organization is developing serverless applications with Oracle Functions. Many functions will need to store state data in a database, which will require using appropriate credentials. However, your corporate security standards mandate encryption of secret information, such as database passwords. How would you address this security requirement?

A. Use OCI Console to enter the password in the function configuration section in the provided input field.

B. Leverage application-level configuration variables to store passwords because they are automatically encrypted by Oracle Functions.

C. Use the OCI Vault service to auto-encrypt the password and then set an application-level configuration variable to reference the auto-decrypted password inside your function container.

D. Encrypt the password using the OCI Vault service and then decrypt this password in your function code with the generated key.

Answer: (SHOW ANSWER)

The best way to store and use secret information, such as database passwords, in Oracle Functions is to use the OCI Vault service. The OCI Vault service provides encryption and decryption capabilities for sensitive data. You can use the OCI Vault service to encrypt the password and store it as an application-level configuration variable. Then, you can use the generated key to decrypt the password in your function code when you need to access the database. Verified References: Oracle Functions: Using Key Management To Encrypt And Decrypt Configuration Variables

NEW QUESTION: 54
What is the maximum allowable size for a secret bundle stored in the Oracle Cloud Infrastructure (OCI) Vault service?
A. 10 KB
B. 25 KB
C. 50 KB
D. 100 KB
Answer: (SHOW ANSWER)
Oracle Secret Management documentation explicitly specifies that themaximum allowable size for a secret bundle is 25 KB. This limit applies when creating a secret or creating a new version of an existing secret.
When using the OCI Console, administrators can enter plain-text secret contents and the Console performs the required Base64 encoding. When communicating directly with the service through supported APIs or CLI operations, secret contents must be supplied in the required encoded form. Secrets are intended for security- sensitive values such as passwords, API tokens, certificates, and configuration credentials rather than arbitrary large files. Payloads exceeding the supported secret-bundle size should therefore be redesigned, divided appropriately, or stored using a service suited to larger objects. Consequently, B is the verified Oracle limit.

NEW QUESTION: 55
Your company requires strict compartment isolation for an event-driven automation framework. You have deployed an Oracle Functions function inside a secure compartment named SecOps-Fn-Compartment. An OCI Events rule has been configured in a separate compartment named App-Resources-Compartment to capture critical resource state changes. Which configuration approach is required to allow the OCI Events service to successfully invoke your function?
A. Define a dynamic group for the rule resource and allow it to manage functions in the target compartment.
B. Define an IAM policy allowing the service principal cloudevents to use functions in the function ' s compartment.
C. Define an IAM policy allowing the service principal oraclefunctions to inspect events in the resource compartment.
D. Define an IAM policy allowing the user group executing the action to manage functions across both compartments.
Answer: (SHOW ANSWER)
OCI Events invokes configured action resources as a managed service rather than by assuming the identity of the human user who created the rule. Oracle ' s Events material documents the service principal cloudEvents and the policy pattern allowing that service to use functions-family. Scoping that permission to the compartment containing the target Function provides the required compartment isolation. A dynamic group is inappropriate because an Events rule is not a Compute instance or workload principal. Granting permissions to oraclefunctions reverses the invocation relationship, while granting a user group broad function- management rights does not

authorize the Events runtime itself. Current Oracle documentation additionally represents cross-tenancy delivery using an eventrule principal and FN_INVOCATION; within the source scenario, B is the intended service authorization model.

NEW QUESTION: 56
Which component of the Docker architecture is the persistent background service responsible for managing container lifecycles, local images, networks, and storage volumes?
A. Docker Registry
B. Docker Client
C. Docker Daemon
D. Docker Compose
Answer: (SHOW ANSWER)
The Docker daemon, implemented as the dockerd background process, runs continuously on a Docker host and performs the server-side container-management operations requested through the Docker API. It manages container execution, images, storage drivers and volumes, and networking. The Docker client is the command interface that sends requests to that daemon; it does not itself manage container lifecycles. A Docker registry such as OCIR stores and distributes images remotely, while Docker Compose describes and coordinates multi- container application configurations. Oracle ' s Linux container documentation explicitly refers to the Docker Engine daemon and its dockerd process, and Oracle Functions troubleshooting likewise distinguishes the local Docker client/daemon from the remote registry. Therefore, the persistent management component described in the question is the Docker Daemon.

NEW QUESTION: 57
Which security capability in Oracle Cloud Infrastructure Container Engine for Kubernetes (OKE) enables developers to scope IAM permissions at the individual service account level for application pods rather than the host Compute instances?
A. Instance Principals
B. Resource Principals
C. Workload Identity
D. Kubernetes RBAC
Answer: (SHOW ANSWER)
OKE Workload Identity enables applications running in Kubernetes pods to receive fine-grained OCI IAM permissions based on their Kubernetes identity rather than inheriting broad permissions from the worker node.
Oracle defines a workload identity using the combination of the OKE cluster, Kubernetes namespace, and Kubernetes service account. IAM policies can therefore restrict access to a specific workload using conditions such as request.principal.type= ' workload ' , request.principal.namespace, request.principal.service_account, and request.principal.cluster_id. Instance Principals grant identity to Compute instances, while Kubernetes RBAC controls access to Kubernetes resources rather than OCI services. Resource Principals apply to supported OCI

resources but do not provide OKE ' s service-account-specific pod identity model. Thus, Workload Identity is the correct capability.

NEW QUESTION: 58
Assuming that your function does NOT have the --provisioned-concurrency option enabled, which parameter is used to configure the time period during which an idle function will remain in memory before Oracle Functions removes its container image from memory?
A. timeout
B. access-timeout
C. idle-timeout
D. None, as this time is not configurable.
Answer: (SHOW ANSWER)
Idle-timeout is the parameter that is used to configure the time period during which an idle function will remain in memory before Oracle Functions removes its container image from memory2. The idle-timeout parameter is specified in seconds and can be set when creating or updating a function2. The default value for idle-timeout is 30 seconds and the maximum value is 900 seconds (15 minutes)2. If a function has the -- provisioned-concurrency option enabled, the idle-timeout parameter is ignored and the function instances are always kept in memory3.Verified References: Creating Functions, Provisioned Concurrency

NEW QUESTION: 59
You deployed a Python application to an Oracle Container Engine for Kubernetes (OKE) cluster. However, while testing you found a bug, which you rectified and then created a new Docker image. You now need to ensure that if this new image does not work once deployed, you should be able to roll back to the previous version. Using kubect1, which strategy should you use?
A. Blue/Green Deployment
B. Canary Deployment
C. Rolling Update
D. A/B Testing
Answer: C (LEAVE A REPLY)
A rolling update is a deployment strategy that gradually replaces the old version of an application with the new version without any downtime4. OKE supports rolling updates by using the kubectl rollout command4. A rolling update allows you to roll back to the previous version if something goes wrong with the new version4.
Therefore, using a rolling update strategy with kubectl ensures that you can roll back to the previous version of your Python application if the new image does not work once deployed.Verified References: Deploy Oracle Container Engine for Kubernetes

NEW QUESTION: 60
When developing microservices, each one can be developed in the language of choice. Which term describes this type of development? (Choose the best answer.)
Freecram.net

A. Agile
B. DevOps
C. Distributed
D. Polyglot
Answer: (SHOW ANSWER)
The term that describes developing microservices in different languages of choice is " Polyglot. " In a polyglot architecture, each microservice is developed using the most appropriate programming language or technology stack for its specific requirements. This approach allows developers to leverage the strengths of different languages and frameworks, enabling them to use the most suitable tool for each microservice while still maintaining interoperability between services.

NEW QUESTION: 61
Which statement is NOT valid regarding context variables used for dynamic backend routing in Oracle Cloud Infrastructure API Gateway?
A. The context variable request.headers[< key >] can be used to dynamically select a backend based on the value of a specific HTTP header in the incoming request.
B. The context variable request.auth[< key >] can be used to dynamically select a backend based on claims returned by an authorizer function or contained in a JWT.
C. The context variable request.query[< key >] can be used to dynamically select a backend based on the value of a specific query parameter in the incoming request.
D. The context variable request.body[< key >] can be used to dynamically select a backend based on parsing specific key-value fields from an incoming JSON payload.
Answer: (SHOW ANSWER)
Oracle API Gateway documents the context-variable selectors available for dynamic backend routing.
Supported selectors include request.auth[< key >], request.headers[< key >], request.host, request.path[< key >], request.query[< key >], request.subdomain[< key >], and usage-plan information. Therefore, headers, authentication claims, and query parameters can all participate directly in backend-selection rules.
Although API Gateway can expose the request body to amulti-argument authorizer function, Oracle does not list request.body[< key >] as a supported selector for dynamic backend routing. Its presence in that separate authorizer context does not make it a dynamic-routing selector capable of parsing arbitrary JSON fields.
Consequently, A, B, and C describe supported routing context variables, while D incorrectly extends the routing model to request-body key selection.

Valid 1z0-1084-26 Dumps shared by ExamDiscuss.com for Helping Passing 1z0-1084-26 Exam! ExamDiscuss.com now offer the **newest 1z0-1084-26 exam dumps**, the

ExamDiscuss.com 1z0-1084-26 exam **questions have been updated** and **answers have been corrected** get the **newest** ExamDiscuss.com 1z0-1084-26 dumps with Test Engine here: https://www.examdiscuss.com/Oracle/exam/1z0-1084-26/premium/ (**182** Q&As Dumps, **35%OFF Special Discount Code: freecram**)

NEW QUESTION: 62
You are configuring a deployment script running in an OCI Compute Instance that needs to retrieve database credentials stored in OCI Vault. Which OCI CLI command should your script execute to retrieve the secret bundle containing the credentials?
A. oci vault secret get --secret-id < secret_id >
B. oci secrets secret-bundle get --secret-id < secret_id >
C. oci kms secret retrieve --secret-id < secret_id >
D. oci vault secret-bundle retrieve --secret-id < secret_id >
Answer: (SHOW ANSWER)
Oracle separates secret-management operations from retrieval of secret contents. The current OCI CLI command for retrieving a secret bundle is oci secrets secret-bundle get --secret-id < secret_id > . If no specific version, stage, or version name is supplied, Oracle returns the bundle associated with the secret version marked CURRENT. The secrets data-plane command retrieves the actual secret bundle, including its encoded content and version properties. The alternative command paths shown in A, C, and D do not represent the documented OCI CLI operation for obtaining secret contents. Oracle ' s current Secrets Management documentation provides the same command directly when describing how applications and administrators retrieve a secret ' s contents. Therefore, option B is the exact verified command.

NEW QUESTION: 63
Which statement is NOT valid regarding the configurations and behaviors of Oracle Cloud Infrastructure (OCI) Queue?
A. The maximum retention period of a message is configurable from 10 seconds to 7 days and cannot be changed after the queue has been created.
B. The visibility timeout of a message is configurable from 1 second to 12 hours at the queue level and can be extended by a consumer during processing.
C. The maximum payload size of a single message published to an Oracle Cloud Infrastructure (OCI) Queue can be configured up to 1 MB.
D. The maximum number of delivery attempts before a message is transferred to a dead letter queue is configurable to values from 1 to 20.
Answer: (SHOW ANSWER)
Statement C is NOT valid because OCI Queue currently limits an individual message to256 KB, not a configurable 1 MB. Oracle ' s current functional limits specify a maximum message size of 256 KB, while a PutMessage batch request can contain up to 20 messages with a combined request size of 512 KB. Message retention can be configured from 10 seconds through seven days and cannot be changed after queue creation.

Queue-level visibility timeout ranges from one second to twelve hours, and visibility can be adjusted while processing. Oracle also documents maximum delivery attempts as configurable from 1 through 20 before failed messages are moved to the dead letter queue. Therefore, C alone contradicts the current Queue functional limits.

NEW QUESTION: 64
Which is NOT a valid backend-type option available when configuring an Oracle Cloud Infrastructure (OCI) API Gateway Deployment?
A. HTTP_BACKEND
B. ORACLE STREAMS_BACKEND
C. ORACLE_FUNCTIONS_BACKEND
Answer: (SHOW ANSWER)
When configuring an OCI API Gateway deployment, you need to specify the backend type for each route in your API deployment specification3. The backend type determines how the API gateway handles requests to that route and forwards them to the appropriate backend service3. The following backend types are valid options for an OCI API Gateway deployment3:
* HTTP_BACKEND: The API gateway forwards requests to an HTTP or HTTPS URL as the backend service.
* ORACLE_FUNCTIONS_BACKEND: The API gateway invokes an Oracle Functions function as the backend service.
* STOCK_RESPONSE_BACKEND: The API gateway returns a stock response without invoking any backend service. ORACLE STREAMS_BACKEND is not a valid backend type for an OCI API Gateway deployment. Oracle Streams is a fully managed, scalable, and durable messaging service that you can use to ingest and consume large amounts of data in real-time4. However, Oracle Streams is not supported as a backend service for an OCI API Gateway deployment.

NEW QUESTION: 65
Which is the smalled unit of Kubernetes architecture?
A. Node
B. Container
C. Cluster
D. Pod
Answer: (SHOW ANSWER)
The smallest unit of Kubernetes architecture is a Pod. A Pod is a logical grouping of one or more containers that are deployed together on the same host and share the same network namespace, storage, and other resources. It represents the smallest deployable unit in Kubernetes and is used to encapsulate and manage one or more closely related containers. Containers within a Pod are scheduled and deployed together, allowing them to communicate and share resources efficiently.

NEW QUESTION: 66

A service you are deploying to Oracle Cloud Infrastructure (OCI) Container Engine for Kubernetes (OKE) uses a docker image from a private repository in OCI Registry (OCIR). Which configuration is necessary to provide access to this repository from OKE?

A. Create a docker-registry secret for OCIR with API key credentials on the cluster, and specify the imagePullSecret property in the application deployment manifest.

B. Create a docker-registry secret for OCIR with identity Auth Token on the cluster, and specify the imagePullSecret property in the application deployment manifest.

C. Create a dynamic group for nodes in the cluster, and a policy that allows the dynamic group to read repositories in the same compartment.

D. Add a generic secret on the cluster containing your identity credentials. Then specify a registryCredentials property in the deployment manifest.

Answer: (SHOW ANSWER)

The necessary configuration to provide access to a private repository in OCI Registry (OCIR) from OCI Container Engine for Kubernetes (OKE) is to create a docker-registry secret for OCIR with an identity Auth Token on the cluster and specify the imagePullSecret property in the application deployment manifest. Here ' s the breakdown of the steps: Create a docker-registry secret for OCIR with an identity Auth Token: In order to authenticate with the private repository in OCIR, you need to create a secret in your OKE cluster that contains the necessary credentials. This can be done by generating an identity Auth Token from the OCI Console and creating a secret in the cluster using the kubectl command. Specify the imagePullSecret property in the application deployment manifest: In your application ' s deployment manifest (such as a Kubernetes Deployment or StatefulSet YAML file), you need to include the imagePullSecret property and specify the name of the secret you created in the previous step. This allows the OKE cluster to use the credentials from the secret to pull the docker image from the private repository in OCIR during deployment. By following these steps, you can ensure that your OKE cluster has the necessary access to the private repository in OCIR, and your application can successfully pull the required docker image during deployment.

NEW QUESTION: 67

Your DevOps team wants to configure a critical alert in OCI Monitoring for a payment gateway microservice application. The security architecture mandate requires that an alert must only fire if the average transaction latency exceeds 500 milliseconds over a sliding 15-minute interval. To avoid individual instance alerts, the query must aggregate all resource metrics within the target compartment together into a single unified stream before checking the threshold.

Which Monitoring Query Language (MQL) expression correctly satisfies all these requirements?

A. TransactionLatency[15m].groupBy(resourceId).mean() > 500

B. TransactionLatency[15m].grouping().mean() > 500

C. TransactionLatency[15m].filter(resourceId).mean() > 500

D. TransactionLatency[15m].absent().mean() > 500

Answer: (SHOW ANSWER)

In OCI Monitoring Query Language, grouping() aggregates multiple matching metric streams into a common result rather than retaining separate streams for individual resources. Oracle ' s documentation demonstrates this syntax with expressions such as IopsRead[1m] {...}.grouping().max(). Here, [15m] establishes the metric aggregation interval, .mean() calculates the average transaction latency, and > 500 supplies the alarm threshold. groupBy(resourceId) would preserve separate result streams by resource ID, which directly contradicts the requirement to evaluate all resource metrics as one unified stream. filter(resourceId) is not valid grouping syntax, and absent() is designed to detect missing metric data rather than evaluate latency values. Therefore, TransactionLatency[15m].grouping().mean() > 500 correctly implements the stated aggregation and threshold logic.

NEW QUESTION: 68
Which of the following is NOT a criterion that is usually met by a microservice?
A. Organized around business capabilities.
B. Tightly coupled
C. Highly maintainable
D. Independently deployable
Answer: (SHOW ANSWER)
The correct answer is: " Tightly coupled. " Tightly coupling is not a criterion that is usually met by a microservice. In fact, microservices are designed to be loosely coupled. Loosely coupling refers to reducing dependencies and minimizing the direct interactions between different components or services. Microservices promote independence and autonomy, allowing each service to operate independently without being tightly bound to other services. The other options listed are criteria that are typically met by microservices:
Organized around business capabilities: Microservices architecture suggests designing services around specific business capabilities or functionalities. This allows for focused and specialized services that align with the organization ' s business needs. Independently deployable: Microservices are designed to be independently deployable units. Each microservice can be developed, tested, and deployed separately, without impacting other services. This enables agility and scalability in the deployment process. Highly maintainable:
Microservices are often designed to be highly maintainable. They are smaller in scope and focused on specific tasks, making it easier to manage and maintain individual services. Additionally, microservices can be updated, patched, or replaced without affecting the entire system, facilitating easier maintenance and evolution of the application. Therefore, the criterion that is NOT typically met by a microservice is being tightly coupled.

NEW QUESTION: 69
Which THREE are valid statements regarding the OCI Container Engine for Kubernetes (OKE) service?
(Choose three.)

A. You must have access to an Oracle Cloud Infrastructure tenancy. Your tenancy must have sufficient quota on different types of resources.
B. OKE cannot use existing network resources for the creation of a new cluster.
C. OKE automatically creates and configures new network resources for the new cluster.
D. There is a limit of three clusters within each region, but there is no limit on the number of nodes and pods you can create within each cluster.
Answer: (SHOW ANSWER)
The valid statements regarding the OCI Container Engine for Kubernetes (OKE) service are: OKE automatically creates and configures new network resources for the new cluster. When creating a new OKE cluster, the service automatically provisions and configures the necessary network resources, such as VCNs, subnets, route tables, security lists, and load balancers, to support the cluster. Your tenancy must have sufficient quota on different types of resources. Before creating an OKE cluster, you need to ensure that your Oracle Cloud Infrastructure (OCI) tenancy has sufficient quota for the required resources, such as compute instances, block storage, networking resources, and load balancers. You must have access to an Oracle Cloud Infrastructure tenancy. To use the OKE service, you need to have access to an OCI tenancy. This means you must have a valid OCI account and the necessary permissions to create and manage resources within the tenancy. The following statements are not valid: OKE cannot use existing network resources for the creation of a new cluster. OKE creates new network resources specifically for the cluster, and it does not support using existing network resources. There is a limit of three clusters within each region, but there is no limit on the number of nodes and pods you can create within each cluster. This statement is incorrect. There is no specific limit on the number of clusters you can create within a region in OKE. However, there may be certain limits or quotas on resources that can impact the number of clusters you can create.

NEW QUESTION: 70
Your company has recently deployed a new web application that uses Oracle Functions. Your manager instructs you to implement monitoring metrics to manage your systems more effectively. You know that Oracle Functions automatically monitors functions on your behalf and reports metrics via Oracle Cloud Infrastructure (OCI) Monitoring. Which TWO metrics are collected and made available by this feature? (Choose two.)
Freecram.net
A. Amount of CPU used by a function
B. Length of time a function runs
C. Number of times a function Is removed
D. Amount of RAM used by a function
E. Number of times a function is invoked
Answer: (SHOW ANSWER)
The correct answers are: Amount of RAM used by a function: Oracle Functions collects and reports the amount of memory (RAM) used by a function during its execution. This metric helps in monitoring and optimizing the resource consumption of functions. Length of time a function runs: Oracle Functions captures and provides the duration of function executions. This metric allows you to track the performance and responsiveness of your functions and identify any potential

bottlenecks or delays. These metrics provide valuable insights into the resource utilization and performance of your functions, enabling you to monitor and optimize their behavior in the Oracle Cloud Infrastructure (OCI) environment.

NEW QUESTION: 71
Which " Action Type " option is NOT available in an Oracle Cloud Infrastructure (OCI) Events rule definition?
A. Streaming
B. Email
C. Notifications
D. Functions
Answer: (SHOW ANSWER)
An action is a response that you define for the rule to perform when the filter finds a matching event1. The action type specifies the service that you want to invoke by delivering the event message1. The following action types are available in OCI Events rule definition1:
* Streaming: Send to a stream from Oracle Streaming Service.
* Notifications: Send to an Oracle Notification Service topic.
* Functions: Send to an Oracle Functions Service endpoint. Email is not a valid action type for OCI Events rule definition. To send an email as an action, you need to use the Notifications service and subscribe to a topic with an email protocol2.

NEW QUESTION: 72
You have a containerized application that requires access to an Autonomous Transaction Processing (ATP) Database. Which option is NOT valid when the container is deployed in an OKE cluster? (Choose the best answer.)
A. Use Kubernetes secrets to configure environment variables on the container with ATP instance OCID, and OCI API credentials. Then use the CreateConnection API endpoint from the service runtime.
B. Install the Oracle Cloud Infrastructure Service Broker on the Kubernetes cluster and deploy ServiceInstance and ServiceBinding resources for ATP. Then use the specified binding name as a volume in the application deployment manifest.
C. Create a Kubernetes secret with contents from the instance Wallet files. Use this secret to create a volume mounted to the appropriate path in the application deployment manifest.
D. Enable Oracle REST Data Services for the required schemas and connect via HTTPS.
Answer: (SHOW ANSWER)
The option that is not valid for connecting to an Autonomous Transaction Processing (ATP) Database from a container in Kubernetes is: Install the Oracle Cloud Infrastructure Service Broker on the Kubernetes cluster and deploy ServiceInstance and ServiceBinding resources for ATP. Then use the specified binding name as a volume in the application deployment manifest. The Oracle Cloud Infrastructure Service Broker is not used for connecting to an ATP Database from a container in Kubernetes. The Service Broker is used for provisioning and managing cloud

services directly from Kubernetes. It allows you to create and manage instances of OCI services using Kubernetes resources like ServiceInstance and ServiceBinding. To connect to an ATP Database from a container in Kubernetes, you can use one of the following valid options: Enable Oracle REST Data Services for the required schemas and connect via HTTPS. This involves enabling and configuring Oracle REST Data Services (ORDS) for the schemas in the ATP Database. You can then connect to the ATP Database using RESTful endpoints provided by ORDS. Use Kubernetes secrets to configure environment variables on the container with ATP instance OCID and OCI API credentials. Then use the CreateConnection API endpoint from the service runtime. This approach involves configuring the necessary environment variables on the container to provide the ATP instance OCID and OCI API credentials. The application can then use the OCI SDK or REST API (such as the CreateConnection endpoint) to establish a connection to the ATP Database. Create a Kubernetes secret with contents from the instance Wallet files. Use this secret to create a volume mounted to the appropriate path in the application deployment manifest. This method involves creating a Kubernetes secret that contains the necessary credentials from the ATP Database ' s instance wallet files. The secret can then be mounted as a volume in the application deployment, allowing the application to access the required credentials for connecting to the ATP Database. Both options 1 and 3 provide valid approaches for connecting to an ATP Database from a container in Kubernetes, depending on the specific requirements and preferences of the application.

NEW QUESTION: 73

Oracle Functions monitors all deployed functions and collects and reports various metrics. Which is NOT available when viewing the Application metrics in the Oracle Cloud Infrastructure (OCI) Console?

A. The length of time a function runs for.

B. The number of retries made by the function before failing due to an error.

C. The number of requests to invoke a function that failed due to throttling.

D. The number of requests to invoke a function that failed with an error response.

Answer: (SHOW ANSWER)

The option that is NOT available when viewing the Application metrics in the Oracle Cloud Infrastructure (OCI) Console is: " The number of retries made by the function before failing due to an error. " When viewing the Application metrics in the OCI Console for Oracle Functions, you can typically see metrics related to the performance and usage of your functions. These metrics provide insights into how your functions are performing and being utilized. The following metrics are usually available: The number of requests to invoke a function that failed due to throttling: This metric indicates the number of requests that were not processed by the function due to reaching the configured concurrency limit or throttling settings. The length of time a function runs for: This metric represents the duration of each function invocation, measuring the time it takes for the function to complete its execution. The number of requests to invoke a function that failed with an error response: This metric counts the number of requests that encountered an error during the function invocation, resulting in a failed response. However, the number of retries made by

the function before failing due to an error is not typically available as part of the Application metrics in the OCI Console.
The retries made by the function are usually handled at the invoker level, and the specific details of retries may not be captured as part of the application-level metrics. It ' s important to note that the availability of metrics and their specific details may vary depending on the version and configuration of Oracle Functions and the monitoring setup. It is recommended to refer to the Oracle Functions documentation and consult the official documentation for accurate and up-to-date information on available metrics.

NEW QUESTION: 74
You are tasked with configuring access for a new developer group named ops-team in OCI. Members of this group must be able to publish raw data points to a custom metric namespace named AppPerformance in the tenancy, and they must also be able to run queries to view both service and custom metrics in a target compartment named production-cmp.
Which TWO IAM policy statements must be written to satisfy these requirements?
A. Allow group ops-team to publish metrics in tenancy where target.metrics.namespace= ' AppPerformance '
B. Allow group ops-team to read metrics in compartment production-cmp to query application metrics
C. Allow group ops-team to manage metrics in compartment production-cmp to query application metrics
D. Allow group ops-team to inspect metrics in compartment production-cmp to list application metrics
E. Allow group ops-team to use metrics in tenancy where target.metrics.namespace= ' AppPerformance '
Answer: B,E (LEAVE A REPLY)
OCI Monitoring IAM maps the use metrics verb to metric write capability, including METRIC_WRITE and the PostMetricData operation. Oracle specifically documents the pattern Allow group < group > to use metrics in tenancy where target.metrics.namespace= ' < namespace > ' for publishing custom metric data.
Therefore, E grants publishing access specifically to AppPerformance. Querying metric data requires read metrics, which includes METRIC_READ and SummarizeMetricsData; B consequently gives the group the required query access in production-cmp. publish is not an OCI IAM policy verb, inspect permits listing metric metadata but not retrieving data points, and manage provides unnecessary privilege. B and E therefore implement the stated requirements using the appropriate least-privilege Monitoring verbs.

NEW QUESTION: 75
(CHK_4 > 2) Which TWO statements are NOT valid regarding the Oracle Cloud Infrastructure (OCI) Streaming service? (Choose two.)
A. OCI Streaming stores all data for 24 hours by default, but that can be extended up to 7 days.B

B. Although OCI Streaming automatically encrypts all data while in transit, it is the developer ' s responsibility to encrypt data at rest, if needed.
C. The throughput of a stream is defined by a partition. A partition provides 1 MB/sec data input and 2 MB
/sec data output.
D. A stream can be configured with either a public or a private endpoint with support for customer managed encryption keys.
E. OCI Streaming can support up to 2,000 requests per second to each partition.
Answer: (SHOW ANSWER)
The two statements that are NOT valid regarding the Oracle Cloud Infrastructure (OCI) Streaming service are: A stream can be configured with either a public or a private endpoint with support for customer managed encryption keys. This statement is not valid because the OCI Streaming service currently supports only private endpoints. Customer managed encryption keys are not currently supported for OCI Streaming. OCI Streaming can support up to 2,000 requests per second to each partition. This statement is not valid because the throughput of a stream is not defined by the partition in terms of requests per second. The throughput of a stream is defined in terms of data input and output rates. Each partition provides 1 MB/sec data input and 2 MB/sec data output, but it does not correspond to a specific number of requests per second. The other statements are valid: OCI Streaming stores all data for 24 hours by default, but that can be extended up to 7 days. Although OCI Streaming automatically encrypts all data while in transit, it is the developer ' s responsibility to encrypt data at rest, if needed.

NEW QUESTION: 76
Which of these is a valid use case for OCI Queue?
A. Managing network traffic between services
B. Storing and retrieving large files
C. Sending real-time streaming data
D. Building decoupled and scalable systems
Answer: (SHOW ANSWER)
OCI Queue is a fully managed serverless service that helps decouple systems and enable asynchronous operations2. Queue handles high-volume transactional data that requires independently processed messages without loss or duplication2. A valid use case for OCI Queue is building decoupled and scalable systems, such as event-driven architectures or microservices-based applications2. For example, you can use Queue to decouple your application and build an event-driven architecture. Decoupling ensures that individual application components can scale independently and that you can future-proof your design so that as new application components are built, they can publish or subscribe to the queue2.

Valid 1z0-1084-26 Dumps shared by ExamDiscuss.com for Helping Passing 1z0-1084-26 Exam! ExamDiscuss.com now offer the **newest 1z0-1084-26 exam dumps**, the ExamDiscuss.com 1z0-1084-26 exam **questions have been updated** and **answers have been corrected** get the **newest** ExamDiscuss.com 1z0-1084-26 dumps with Test Engine here: https://www.examdiscuss.com/Oracle/exam/1z0-1084-26/premium/ (**182** Q&As Dumps, **35%OFF Special Discount Code: freecram**)

NEW QUESTION: 77
Which credential type is required to act as the password when authenticating the Docker CLI to Oracle Cloud Infrastructure Registry?
A. Console Password
B. OCI Auth Token
C. API Signing Key
D. Customer Secret Key
Answer: (SHOW ANSWER)
Oracle Container Registry requires an OCI Auth Token when Docker prompts for the password during docker login. Oracle explicitly documents that the user first generates an auth token under the OCI user profile, then supplies the tenancy namespace and username as the Docker username and the auth token as the password. A Console password is intended for interactive OCI Console authentication and is not the Docker Registry credential. API signing keys authenticate signed OCI API requests used by the CLI, SDKs, and related tools, while Customer Secret Keys are associated with Object Storage compatibility APIs. Consequently, the programmatic credential designed specifically to serve as the Docker Registry password is the OCI Auth Token.

NEW QUESTION: 78
You are configuring a route on an Oracle Cloud Infrastructure API Gateway to an HTTP backend service that takes up to four minutes to execute long-running analytical queries. Which timeout configuration should you apply in the API deployment specification to safely prevent a gateway timeout?
A. Set the readTimeoutInSeconds parameter to 120.0 seconds.
B. Set the readTimeoutInSeconds parameter to 240.0 seconds.
C. Set the readTimeoutInSeconds parameter to 360.0 seconds.
D. Set the readTimeoutInSeconds parameter to 480.0 seconds.
Answer: (SHOW ANSWER)
For an HTTP backend, OCI API Gateway provides readTimeoutInSeconds to control how long the gateway waits for the backend response after the connection has been established. A four-minute backend processing window requires a timeout of 240 seconds, so B is sufficient without exceeding the platform limit. Oracle documents a maximum configurable timeout value of300 secondsfor HTTP backend timeout parameters.

Therefore, the 360-second and 480-second choices exceed the supported range and would not be valid deployment values. A 120-second timeout would expire before the four-minute analytical operation can complete. Thus, 240 seconds is both within the API Gateway limit and aligned with the backend ' s maximum execution duration.

Valid 1z0-1084-26 Dumps shared by ExamDiscuss.com for Helping Passing 1z0-1084-26 Exam! ExamDiscuss.com now offer the **newest 1z0-1084-26 exam dumps**, the ExamDiscuss.com 1z0-1084-26 exam **questions have been updated** and **answers have been corrected** get the **newest** ExamDiscuss.com 1z0-1084-26 dumps with Test Engine here: https://www.examdiscuss.com/Oracle/exam/1z0-1084-26/premium/ (**182** Q&As Dumps, **35%OFF Special Discount Code: freecram**)