Oracle.1z0-1084-24.v2025-03-10.q34

Exam Code:	1z0-1084-24
Exam Name:	Oracle Cloud Infrastructure 2024 Developer Professional
Certification Provider:	Oracle
Free Question Number:	34
Version:	v2025-03-10
# of views:	491
# of Questions views:	340
https://www.freepdfdumps.com/Oracle.1z0-1084-24.v2025-03-10.q34.html	

NEW QUESTION: 1
As a Cloud Native developer, you have written a web service for your company. However, your security team has suggested that your web service should address Distributed Denial-of-Service (DDoS) attack. You are time-constrained and you need to ensure that this is implemented as soon as possible. What should you do in this scenario? (Choose the best answer.)
A. Use a third party service integration to Implement DDoS attack mitigation.
B. Re-write your web service and implement rate limiting.
C. Use the OCI Virtual Cloud Network (VCN) segregation to control DDoS.
D. Use the OCI API Gateway service and configure rate limiting.
Answer: (SHOW ANSWER)
The correct answer in this scenario is to use the OCI API Gateway service and configure rate limiting. Using the OCI API Gateway service and configuring rate limiting is an effective approach to address Distributed Denial-of-Service (DDoS) attacks. By implementing rate limiting, you can control the number of requests that can be made to your web service within a specific time frame. This helps to prevent overload and ensures that your service can handle legitimate traffic while mitigating the impact of DDoS attacks. By leveraging the OCI API Gateway service, you can easily configure rate limiting rules to restrict the number of requests per second or per minute. This allows you to set appropriate thresholds and safeguard your web service from being overwhelmed by excessive requests. The API Gateway acts as a protective layer, filtering out malicious traffic and ensuring the smooth operation of your service. While options like OCI Virtual Cloud Network (VCN) segregation and third-party service integrations may contribute to overall security, they do not specifically address DDoS attacks as efficiently as rate limiting. VCN segregation focuses more on network segmentation and isolation, while third-party service integration may introduce additional dependencies and complexities.

Re-writing your web service and implementing rate limiting is a viable option, but it may not be feasible considering the time constraints mentioned. Leveraging the OCI API Gateway service provides a quicker and easier solution to implement DDoS attack mitigation through rate limiting.

NEW QUESTION: 2
Which testing strategy achieves high velocity of deployments and releases of cloud native applications?
(Choose the best answer.)
A. Penetration testing
B. Automated testing
C. Integration testing
D. A/B testing
Answer: B (LEAVE A REPLY)
The testing strategy that achieves high velocity of deployments and releases of cloud native applications is
"Automated testing." Automated testing involves the use of automated tools and frameworks to execute tests, validate functionality, and detect issues or bugs in an application. By automating the testing process, developers and DevOps teams can rapidly test and validate code changes, ensuring that new features and updates are functioning correctly before being deployed to production. This approach helps increase the speed and efficiency of the testing process, allowing for faster and more frequent deployments of cloud native applications.

NEW QUESTION: 3
Your organization has mandated that all deployed container images used for microservices must be signed by a specified master encryption key (MEK). You have appropriately signed the container images as part of your build process, but must now ensure that they are automatically verified when they are deployed to Oracle Cloud Infrastructure (OCI) Container Engine for Kubemetes (OKE) clusters. Which option should be used to mandate image verification when deploying to OKE clusters, assuming that MEK is already stored in an available OCI Vault? (Choose the best answer.)
A. Enable image verification policies separately for each Kubemetes pod deployment because this is enforced at the pod level.
B. Enable image verification policies separately for each node pool within each OKE cluster because this is enforced at the node pool level.
C. Enable image verification policies separately for each OKE cluster because this is enforced at the cluster level.
(Correct)
D. Enable Image verification policies for your OKE service control plane which will enforce this for all OKE clusters.

Answer: (SHOW ANSWER)
To mandate image verification when deploying container images to Oracle Cloud Infrastructure (OCI) Container Engine for Kubernetes (OKE) clusters, you should enable image verification policies separately for each OKE cluster. This is enforced at the cluster level. Enabling image verification policies at the cluster level ensures that all container images deployed to the OKE cluster are automatically verified against the specified master encryption key (MEK). This helps maintain the security and integrity of the deployed microservices by ensuring that only signed and trusted container images are used. Enabling image verification policies at the cluster level allows for consistent and centralized enforcement of the verification process across all nodes and node pools within the cluster. It provides a standardized approach to image verification for the entire cluster, simplifying management and ensuring compliance with the organization's mandate. Enabling image verification policies separately for each node pool or at the pod level would introduce complexity and potential inconsistencies in the verification process. Therefore, enforcing image verification at the cluster level is the recommended approach.

NEW QUESTION: 4
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: A (LEAVE A REPLY)
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: 5
You have just finished building and compiling the software required to implement the API microservice component. You need to rebuild the API docker image, and plan to tag it as: ocIdevops/api:latest Which docker command would re-create the API docker image?
A. docker build -t OCIdevops/api:latest
B. docker create -t OCIdevops/api:latest
C. docker image -t OCIdevops/api:latest

D. docker compile -t OCI devops/api:latest

Answer: A (LEAVE A REPLY)

The correct command to rebuild the API docker image and tag it as OCIdevops/api:latest is: docker build -t OCIdevops/api:latest The docker build command is used to build a Docker image from a Dockerfile. The -t flag is used to specify the name and optionally a tag for the image. In this case, the name of the image is OCIdevops/api and the tag is latest. By running this command, the Docker image will be recreated based on the instructions in the Dockerfile and tagged with the specified name and tag.

NEW QUESTION: 6

Which two "Action Type" options are NOT available in an Oracle Cloud Infrastructure (OCI) Events rule definition? (Choose two.)

A. Email

B. Streaming

C. Slack

D. Functions

E. Notifications

Answer: A,C (LEAVE A REPLY)

The two "Action Type" options that are NOT available in an Oracle Cloud Infrastructure (OCI) Events rule definition are: Email (Correct) Slack (Correct) The available "Action Type" options in OCI Events rule definition include Functions, Notifications, and Streaming. However, email and Slack are not directly supported as action types in OCI Events. Instead, you can use Notifications to send notifications to various notification channels, including email and Slack, through the OCI Notifications service.

NEW QUESTION: 7

Which testing measure should be considered when using test cases that simultaneously validate a deployment and perform a selected set of functional tasks?

A. Resource Utilization

B. Functionality

C. Scalability

D. Robust Deployment

E. Resiliency

Answer: D (LEAVE A REPLY)

The correct answer is: "Robust Deployment." When using test cases that simultaneously validate a deployment and perform a selected set of functional tasks, the testing measure that should be considered is

"Robust Deployment." Robust Deployment refers to the ability of an application or system to be deployed reliably and consistently, without errors or failures. It involves ensuring that the deployment process is well- defined, automated, and able to handle different scenarios and configurations. When conducting testing that combines the validation of deployment

and functional tasks, it is crucial to ensure that the deployment itself is robust. This means verifying that the application or system can be successfully deployed and configured without encountering deployment-related issues such as incorrect configurations, missing dependencies, or compatibility problems. By considering "Robust Deployment" as a testing measure, you can evaluate the reliability and effectiveness of the deployment process, ensuring that the application or system is deployed correctly and ready to perform the selected set of functional tasks.

NEW QUESTION: 8
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: A,E (LEAVE A REPLY)
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: 9
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: 10
Which command is used to get a Docker image from Oracle Cloud Infrastructure Registry (OCIR) to the client machine?
A. docker pull <region-key>.ocir.io/<tenancy-namespace>/<repo-name>: <tag>
B. docker pull <tenancy-namespace>/<region-key>.ocir.io/<repo-name>: <tag>
C. docker fetch <region-key>.ocir.io/<tenancy-namespace>/<repo-name>:<tag>
D. docker fetch <tenancy-namespace>/<region-key>.ocir.io/<repo-name>:<tag>
Answer: (SHOW ANSWER)
To pull a Docker image from OCI Registry to the client machine, you need to use the docker pull command with the following syntax1: docker pull <region-key>.ocir.io/<tenancy-namespace>/<repo-name>:<tag> where:
* <region-key> is the key for the OCI Registry region you're using. For example, iad. See Availability by Region1.
* ocir.io is the OCI Registry name.
* <tenancy-namespace> is the auto-generated Object Storage namespace string of the tenancy that owns the repository from which you want to pull the image (as shown on the Tenancy Information page)1.
* <repo-name> is the name of the repository that contains the image you want to pull.
* <tag> is the tag of the image you want to pull.

NEW QUESTION: 11

You are developing a serverless application with Oracle Functions and Oracle Cloud Infrastructure Object Storage. Your function needs to read a JSON file object from an Object Storage bucket named "input-bucket" in compartment "qa-compartment". Your corporate security standards mandate the use of Resource Principals for this use case. Which two statements are needed to implement this use case? (Choose two.)
A. Set up a policy to grant all functions read access to the bucket: allow all functions in compartment qa- compartment to read objects in target.bucket.name= "input-bucket'
B. Set up a policy to grant your user account read access to the bucket: allow user XYZ to read objects in compartment qa-compartment where target.bucket.name= "input-bucket'
C. Set up the following dynamic group for your function's OCID: Name: read-file-dg Rule: resource.id = "ocid1.fnfunc.oc1.phx.aaaaaaaakeaobctakezjz5i4ujj7g25q7sx5m vr55pms6f4da'
D. No policies are needed. By default, every function has read access to Object Storage buckets in the tenancy.
E. Set up a policy with the following statement to grant read access to the bucket: allow dynamic-group read-file-dg to read objects in compartment qa- compartment where target.bucket.name= 'input-bucket'
Answer: C,E (LEAVE A REPLY)
The correct answers are: Set up the following dynamic group for your function's OCID: Name: read-file-dg Rule: resource.id = "ocid1.fnfunc.oc1.phx.aaaaaaaakeaobctakezjz5i4ujj7g25q7sx5mvr55pms6f4da" Set up a policy with the following statement to grant read access to the bucket: Statement: allow dynamic-group read- file-dg to read objects in compartment qa-compartment where target.bucket.name = 'input-bucket' Explanation: To implement the use case of reading a JSON file object from an Object Storage bucket using Resource Principals with Oracle Functions, you need to configure the following: Create a dynamic group named "read-file-dg" and associate it with your function's OCID. This dynamic group helps identify the function as a member of the group for policy enforcement. Create a policy that grants read access to the bucket. The policy statement should allow the dynamic group "read-file-dg" to read objects in the compartment "qa-compartment" and specify the target bucket name as "input-bucket". This policy ensures that the function has the necessary permissions to access the specified bucket. By setting up the dynamic group and policy, you ensure that the function, as a member of the dynamic group, has the required read access to the specified Object Storage bucket in the specified compartment.

NEW QUESTION: 12
Which feature is typically NOT associated with Cloud Native?
A. Immutable Infrastructure
B. Declarative APIs
C. Containers
D. Application Servers

E. Service Meshes

Answer: D (LEAVE A REPLY)

The feature that is typically NOT associated with Cloud Native is "Application Servers." Cloud Native architecture emphasizes lightweight, scalable, and containerized deployments, which often replace traditional monolithic application servers. Instead of relying on application servers, Cloud Native applications are typically deployed as containerized microservices that can be orchestrated and managed using container orchestration platforms like Kubernetes. This approach enables greater flexibility, scalability, and agility in deploying and managing applications. While application servers have been widely used in traditional application architectures, they are not a characteristic feature of Cloud Native architectures. Cloud Native architectures focus on containerization, declarative APIs, immutable infrastructure, and service meshes to enable efficient and scalable deployment and management of applications.

NEW QUESTION: 13

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.)

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: A,D (LEAVE A REPLY)

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: 14

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: C (LEAVE A REPLY)
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: 15
Your organization has deployed their e-commerce application on Oracle Container Engine for Kubernetes (OKE) and they are using the Oracle Cloud Infrastructure Registry (OCIR) service as their Docker image repository. They have deployed the OKE cluster using the 'custom create' option, and their Virtual Cloud Network (VCN) has three public subnets with associated Route Tables, Security Lists, and Internet Gateway.
However, their application containers are failing to deploy. On investigation, they discover that the images are not being pulled from the designated OCIR repository, even though the YAML configuration has the correct path to the images. What is a valid concern here that needs to be further investigated?
A. Security List rule for TCP port 22 needs to be added to connect to the OCIR service.
B. VCN hosting the OKE cluster worker nodes needs to have a NAT gateway to access OCIR repositories.
C. Identity and Access Management (IAM) credentials need to be added for each user that deploys applications to the OKE cluster.
D. OKE cluster needs to have a secret with the credentials of their OCIR repository and use that secret in the Kubernetes deployment manifest.

Answer: D (LEAVE A REPLY)
A valid concern that needs to be further investigated in this scenario is whether the OKE cluster has a secret with the credentials of the Oracle Cloud Infrastructure Registry (OCIR) repository and if that secret is being used in the Kubernetes deployment manifest. Here's why this concern is relevant: Access to the OCIR repository: In order for the OKE cluster to pull images from the OCIR repository, it needs proper authentication credentials. These credentials are typically provided in the form of a secret, which contains the necessary information to authenticate with the registry. Secret in the deployment manifest: The Kubernetes deployment manifest defines how the application containers should be deployed. It includes specifications such as the container image, resource requirements, and environment variables. To pull images from a private repository like OCIR, the deployment manifest needs to reference the appropriate secret that contains the registry credentials. If the images are not being pulled from the designated OCIR repository, it suggests that either the secret with the OCIR credentials is missing or it is not properly referenced in the deployment manifest. Further investigation should focus on verifying the presence and correctness of the secret, as well as confirming that it is correctly referenced in the deployment manifest for the application containers. By ensuring the presence of the secret and proper configuration in the deployment manifest, the OKE cluster will have the necessary credentials to access the OCIR repository and successfully deploy the application containers.

NEW QUESTION: 16
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: A (LEAVE A REPLY)
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: 17
How are cloud native application versions deployed to an OKE cluster when using a blue/green deployment strategy?
A. Current applications are slowly replaced with new application versions.
B. New application versions are deployed in minor increments to a select group of people.
C. Both old and new application versions are deployed to production at the same time.
Answer: C (LEAVE A REPLY)
Blue/Green deployment strategy allows releasing a new version of an application using two identical environments where one of them is active at a given time. The current version of the application is provisioned on the active environment, whereas the new version gets deployed to the standby environment1. The traffic is shifted from the active to the standby environment by updating the ingress resource2. Therefore, both old and new application versions are deployed to production at the same time, but only one of them receives the traffic. Verified References: Announcing new deployment strategies for OCI DevOps Service, Blue-Green OKE Deployment

NEW QUESTION: 18
A developer using Oracle Cloud Infrastructure (OCI) API Gateway needs to authenticate the API requests to their web application. The authentication process must be implemented using a custom scheme which accepts string-based parameters from the API caller. Which approach should the developer use in this scenario?
A. Create a cross account functions authorizer.
B. Create an authorizer function using OCI Identity and Access Management 91AM) based authentication.
C. Create an authorizer function using request header authorization.
D. Create an authorizer function using token-based authorization.
Answer: D (LEAVE A REPLY)

In the given scenario, the developer should use the approach of creating an authorizer function using token- based authorization. Token-based authorization is a commonly used approach for authenticating API requests.
It involves generating and issuing tokens to API callers, which they can then include in the requests they make to the API. The tokens serve as proof of authentication and are validated by the server to ensure the caller's identity and access rights. By creating an authorizer function using token-based authorization, the developer can implement a custom scheme that accepts string-based parameters from the API caller. This allows the developer to define their own authentication logic and validate the provided tokens according to their requirements. The authorizer function can be configured in the OCI API Gateway to be invoked before forwarding the request to the web application. It will perform the necessary token validation and authentication checks, allowing only authorized requests to access the protected resources of the web application.

NEW QUESTION: 19
To effectively test your cloud native applications for "unknown unknowns", you need to employ various testing and deployment strategies. Which strategy involves exposing new functionality or features to only a small set of users?
A. A/B Testing
B. Component Testing
C. Blue/Green Deployment
D. Canary Deployment
Answer: (SHOW ANSWER)
The strategy that involves exposing new functionality or features to only a small set of users is called Canary Deployment. Canary deployment is a technique used in software development and deployment where a new version of an application or feature is released to a small subset of users or a specific group of servers. This allows for testing and gathering feedback on the new functionality in a controlled and limited environment before making it available to a wider audience. In a canary deployment, a small portion of the traffic is routed to the new version while the majority of the traffic still goes to the stable version. This allows for monitoring and evaluation of the new functionality in real-world conditions while minimizing the impact of any potential issues or bugs. If the new version performs well and meets the desired criteria, it can then be gradually rolled out to a larger user base or all servers. By exposing the new functionality or features to a small set of users initially, canary deployment helps in identifying any unforeseen issues, gathering feedback, and ensuring the stability and reliability of the application before a full deployment.

NEW QUESTION: 20

You are creating an API deployment in Oracle Cloud Infrastructure (OCI) API Gateway and you want to configure request policies to control access. Which is NOT available in OCI API Gateway?

A. Controlling access to the backend OCI resources.

B. Limiting the number of requests sent to the backend services.

C. Enabling Cross-Origin Resource Sharing (CORS) support.

D. Providing authentication and authorization.

Answer: A (LEAVE A REPLY)

The correct answer is: Controlling access to the backend OCI resources. OCI API Gateway does not provide direct control over access to backend OCI resources. It primarily focuses on managing and securing access to APIs exposed through the gateway. The gateway acts as a front-end for APIs and provides features such as authentication, authorization, rate limiting, and CORS support. While you can configure authentication and authorization policies, limit the number of requests, and enable CORS support in OCI API Gateway, it does not directly control access to backend OCI resources. Access to backend resources is typically managed through other means, such as IAM policies, network security rules, or resource-specific access controls.

NEW QUESTION: 21

(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: D,E (LEAVE A REPLY)

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: 22
To enforce mutual TLS (mTLS) authentication for clients of your microservices, your team has chosen to leverage the Oracle Cloud Infrastructure (OCI) API Gateway service to create new API Deployments that will direct requests to your microservices. Which is NOT valid regarding the mTLS options in OCI API Gateway?
A. Custom CA or custom CA bundles can be added to your gateway's trust store ONLY if they already exist in the OCI Certificates service.
B. Adding a custom certificate authority (CA) or custom CA bundle to your gateway's trust store for mTLS is optional unless you need to reject certificates that do not contain particular values (such as a domain name).
C. The mTLS request policy can only be enabled at the API deployment specification level, which then applies globally to ALL routes in that deployment.
D. Once the mTLS request policy is enabled, ALL requests with valid certificates are routed to the backend unless you have defined one or more particular values (such as a domain name).
Answer: B (LEAVE A REPLY)
The correct answer is: "Adding a custom certificate authority (CA) or custom CA bundle to your gateway's trust store for mTLS is optional unless you need to reject certificates that do not contain particular values (such as a domain name)." The statement that is NOT valid regarding the mTLS options in OCI API Gateway is: "Adding a custom certificate authority (CA) or custom CA bundle to your gateway's trust store for mTLS is optional unless you need to reject certificates that do not contain particular values (such as a domain name)." In OCI API Gateway, adding a custom certificate authority (CA) or custom CA bundle to the gateway's trust store is not optional. It is a necessary step when configuring mTLS authentication. The trust store in the gateway is used to validate the client certificates presented during mTLS authentication. The other options listed are valid regarding the mTLS options in OCI API Gateway: Once the mTLS request policy is enabled, all requests with valid certificates are routed to the backend unless specific values (such as a domain name) are defined. This means that only requests with valid client certificates will be allowed to access the backend microservices. The mTLS request policy can only be enabled at the API deployment specification level, and it applies globally to all routes in that deployment. This ensures consistent mTLS authentication across all routes and endpoints in the API deployment. Custom CA or custom CA bundles can be added to the gateway's trust store, but only if they already exist in the OCI Certificates service. This allows you to include trusted CAs or CA bundles to validate client certificates during mTLS authentication.

NEW QUESTION: 23
You are using Oracle Cloud Infrastructure (OCI) Resource Manager to manage your infrastructure lifecycle and wish to receive an email each time a Terraform action begins. How should you use the OCI Events service to do this without writing any code?
A. Create a rule in OCI Events service matching the "Resource Manager Stack - Update" condition. Then select "Action Type: Email" and provide the destination email address.
B. Create an OCI Notification topic and email subscription with the destination email address. Then create an OCI Events rule matching "Resource Manager Job - Create" condition, and select the notification topic for the corresponding action.
C. Create an OCI Email Delivery configuration with the destination email address. Then create an OCI Events rule matching "Resource Manager Job - Create" condition, and select the email configuration for the corresponding action.
D. Create an OCI Notifications topic and email subscription with the destination email address. Then create an OCI Events rule matching "Resource Manager Stack - Update" condition, and select the notification topic for the corresponding action.
Answer: B (LEAVE A REPLY)
The correct approach to receive an email each time a Terraform action begins in Oracle Cloud Infrastructure (OCI) Resource Manager without writing any code is as follows: Create an OCI Notification topic and email subscription with the destination email address. This will define the email delivery configuration. Create an OCI Events rule that matches the "Resource Manager Job - Create" condition. This rule will be triggered when a Resource Manager job is created. In the OCI Events rule, select the notification topic that was created in step 1 as the action for the corresponding event. This will ensure that the notification is sent to the specified email address. By following these steps, you can configure the OCI Events service to send an email notification whenever a Resource Manager job is created in OCI Resource Manager.

NEW QUESTION: 24
You are building a cloud native serverless travel application with multiple Oracle Functions in Java, Python, and Node.js. You need to build and deploy these functions to a single application named travel-app. Which command will help you complete this task successfully?
A. fn function deploy app travel-app--all
B. fn app deploy --app travel-app --all
C. fn app --app travel-app deploy --ext java pyljs
D. fn deploy--app travel-app --all
Answer: (SHOW ANSWER)
The correct answer is: fn deploy --app travel-app --all Explanation: To build and deploy multiple Oracle Functions as part of a single application named "travel-app," you can use the fn deploy command with the appropriate options. The command fn deploy --app travel-

app --all is the correct syntax. Here's what each part of the command does: fn deploy: This command is used to deploy functions and applications in Oracle Functions. --app travel-app: This option specifies the application name as "travel-app," indicating that you want to deploy functions to this application. --all: This option indicates that you want to deploy all the functions within the application. By using fn deploy --app travel-app --all, you can build and deploy all the functions in your travel application across different programming languages (Java, Python, and Node.js) to the "travel-app" application in Oracle Functions.

NEW QUESTION: 25
Your team has chosen to use master encryption key (MEK) within an Oracle Cloud Infrastructure (OCI) Vault for encrypting Kubernetes secrets associated with your microservice deployments in OCI Container Engine for Kubernetes (OKE) clusters so that you can easily manage key rotation. Which of the following is NOT valid about rotating keys in the OCI Vault service?
A. Once rotated, older key versions can be used for encryption until they are deleted.
B. Both software and HSM-protected MEKS can be rotated.
C. When you rotate an MEK, a new key version is automatically generated.
D. Each key version is tracked internally with separate unique OCIDS.
Answer: A (LEAVE A REPLY)
The correct answer is: "Once rotated, older key versions can be used for encryption until they are deleted." The statement that is NOT valid about rotating keys in the OCI Vault service is: "Once rotated, older key versions can be used for encryption until they are deleted." In the OCI Vault service, when you rotate a master encryption key (MEK), a new key version is automatically generated. However, once a key is rotated and a new version is created, the older key versions are no longer usable for encryption. The purpose of key rotation is to ensure that the encryption keys are regularly updated and that older keys are no longer used to protect sensitive data. This enhances security by minimizing the impact of potential key compromises. The other statements mentioned are valid: Both software and hardware security module (HSM)-protected MEKs can be rotated. This provides flexibility in choosing the type of MEK and ensures that key rotation can be performed regardless of the encryption method used. Each key version is tracked internally with separate unique OCIDs (Oracle Cloud Identifiers). This allows for easy management and tracking of different key versions within the OCI Vault service. In summary, the statement that is NOT valid is the one suggesting that older key versions can still be used for encryption until they are deleted. Key rotation is designed to ensure the use of the latest key version and to retire older key versions to enhance security.

NEW QUESTION: 26
Which TWO are required to access the Oracle Cloud Infrastructure (OCI) Container Engine for Kubernetes (OKE) cluster from the kubectl CLI? (Choose two.)
A. Tiller enabled on the OKE cluster.

B. An SSH key pair with the public key added to the cluster worker nodes.
C. Install and configure the OCI CLI.
D. A configured OCI API signing key pair.
E. OCI Identity and Access Management (IAM) Auth Token.
Answer: C,D (LEAVE A REPLY)
The correct options are: A configured OCI API signing key pair: The API signing key pair is used for authentication and authorization to access OCI resources, including the OKE cluster. The private key should be configured on your local machine to authenticate API requests. An SSH key pair with the public key added to the cluster worker nodes: This is required for secure SSH access to the worker nodes in the OKE cluster.
You need to generate an SSH key pair and add the public key to the cluster's worker node pool during cluster creation or update. Therefore, the correct options are having a configured OCI API signing key pair and an SSH key pair with the public key added to the cluster worker nodes.

NEW QUESTION: 27
(CHK_1>3) You have an e-commerce application that loads customers' transactional data into the Oracle Cloud Infrastructure (OCI) Streaming service. The data must now be extracted and transformed before sending it to a third-party REST endpoint. You have been directed to leverage the OCI Service Connector Hub to automate this process. Which configuration option would address this requirement?
A. Configure a new service connector as follows: * Source: Streaming * Task: Functions * Target:
Functions
B. Configure a new service connector as follows: * Source: Streaming * Task: API Gateway * Target:
Notifications
C. Configure a new service connector as follows: * Source: Streaming * Task: None * Target: Notifications
D. Configure a new service connector as follows: * Source: Streaming * Task: API Gateway * Target:
Functions
E. Configure a new service connector as follows: * Source: Streaming * Task: Functions * Target: API Gateway
Answer: (SHOW ANSWER)
To address the requirement of extracting and transforming data from the Oracle Cloud Infrastructure (OCI) Streaming service and sending it to a third-party REST endpoint using the OCI Service Connector Hub, the best configuration option is: Configure a new service connector as follows: * Source: Streaming * Task: None
* Target: Notifications By selecting the Streaming service as the source, you can capture the transactional data from the stream. Since there is a need to transform and send the

data to a third-party REST endpoint, you don't need to specify any specific task in the connector. The target is set to Notifications, which allows you to send the transformed data to an endpoint outside of the OCI environment. Notifications can be configured to deliver the data to various supported destinations, including HTTP endpoints, email addresses, and more. This configuration enables you to automate the process of extracting data from the streaming service and sending it to the desired third-party REST endpoint, fulfilling the requirement of extracting, transforming, and forwarding the data.

NEW QUESTION: 28
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: A,C,D (LEAVE A REPLY)
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: 29
Which is NOT a valid option to execute a function deployed in Oracle Functions?
A. Invoke from the Docker CLI.

B. Send signed HTTP requests to the function's invoke endpoint.
C. Invoke from the Fn Project CLI.
D. Trigger by an event in the Oracle Cloud Infrastructure (OCI) Events service.
E. Invoke from the OCI CLI.
Answer: A (LEAVE A REPLY)
The correct answer is: Invoke from the Docker CLI. Explanation: Executing a function deployed in Oracle Functions is typically done using the following options: Invoke from the Fn Project CLI: The Fn Project CLI provides a command-line interface specifically designed for interacting with Oracle Functions. You can use commands like fn invoke to invoke a function. Trigger by an event in the Oracle Cloud Infrastructure (OCI) Events service: You can configure events in OCI to trigger your function based on various criteria, such as object storage events, resource state changes, or scheduled events. Invoke from the OCI CLI: The OCI CLI (Command Line Interface) allows you to interact with various services in Oracle Cloud Infrastructure, including Oracle Functions. You can use the fn invoke command to invoke a function. Send signed HTTP requests to the function's invoke endpoint: Oracle Functions provides an HTTP endpoint that can be used to invoke functions. You can send signed HTTP requests to this endpoint using tools or programming languages that support making HTTP requests. On the other hand, invoking a function deployed in Oracle Functions using the Docker CLI is not a valid option. The Docker CLI is primarily used for managing Docker containers and images, and it does not provide a direct mechanism for invoking functions in Oracle Functions.

NEW QUESTION: 30
A company is developing a new application that needs to process transactions in real time. The company wants to ensure that all transactions are processed in order and that no transaction is lost. Which of these is a correct strategy for leveraging OCI Queue in this scenario?
A. Use a separate queue for each type of transaction.
B. Use a single queue to process all transactions.
C. Use a separate queue for each application instance.
D. Use a priority queue to prioritize requests.
Answer: B (LEAVE A REPLY)
OCI Queue is a service for enabling asynchronous (decoupled) communication in a serverless manner3. Queue handles high-volume transactional data that requires independent processing without loss or duplication3. Queue supports ordering of messages within a queue by using the FIFO (first-in-first-out) delivery option3. Therefore, using a single queue to process all transactions ensures that all transactions are processed in order and that no transaction is lost. Verified References: Overview of Queue

NEW QUESTION: 31

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: B (LEAVE A REPLY)
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: 32
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: D (LEAVE A REPLY)
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: 33
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: 34
Which TWO statements are true for serverless computing and serverless architectures? (Choose two.)
A. Serverless function execution is fully managed by third party.
B. Applications running on a FaaS (Functions as a Service) platform.
C. Long running tasks are perfectly suited for serverless.
D. Application DevOps team is responsible for scaling.
E. Serverless function state should never be stored externally.
Answer: A,B (LEAVE A REPLY)
The two true statements for serverless computing and serverless architectures are: Applications running on a FaaS (Functions as a Service) platform: Serverless architectures typically involve running code in the form of functions on a serverless platform. These functions are event-driven and executed in response to specific triggers or events. Serverless function execution is fully managed by a third party: In serverless computing, the cloud provider takes care of the infrastructure management and resource provisioning. The execution of serverless functions is handled automatically by the platform, relieving developers from the responsibility of managing servers or infrastructure. It's important to note that long running tasks are not typically suited for serverless architectures due to the event-driven nature of serverless functions. Also, while serverless functions may have state, it is recommended to avoid external storage dependencies and instead leverage stateless functions whenever possible. Additionally, scaling in serverless architectures is typically handled automatically by the platform, rather than being the responsibility of the application DevOps team.