Oracle.1Z0-1084-21.v2022-03-23.q24

Exam Code:	1Z0-1084-21
Exam Name:	Oracle Cloud Infrastructure Developer 2021 Associate
Certification Provider:	Oracle
Free Question Number:	24
Version:	v2022-03-23
# of views:	815
# of Questions views:	240
https://www.freeqas.com/qa/Oracle/1Z0-1084-21/Oracle.1Z0-1084-21.v2022-03-23.q24.html	

NEW QUESTION: 1
You created a pod called "nginx" and its state is set to Pending.
Which command can you run to see the reason why the "nginx" pod is in the pending state?
A. kubect2 logs pod nginx
B. kubect2 describe pod nginx
C. kubect2 get pod nginx
D. Through the Oracle Cloud Infrastructure Console
Answer: B (LEAVE A REPLY)
Explanation
Debugging Pods
The first step in debugging a pod is taking a look at it. Check the current state of the pod and recent events with the following command:
kubectl describe pods ${POD_NAME}
Look at the state of the containers in the pod. Are they all Running? Have therebeen recent restarts?
Continue debugging depending on the state of the pods.
My pod stays pending
If a pod is stuck in Pending it means that it can not be scheduled onto a node. Generally this is because there are insufficient resources of one type oranother that prevent scheduling. Look at the output of the kubectl describe ... command above. There should be messages from the scheduler about why it can not schedule your pod.
https://kubernetes.io/docs/tasks/debug-application-cluster/debug-pod-replication-controller/

NEW QUESTION: 2
What is the minimum of storage that a persistent volume claim can obtain in Oracle Cloud Infrastructure Container Engine for Kubernetes (OKE)?
A. 50 GB

B. 10 GB
C. 1 GB
D. 1 TB
Answer: A (LEAVE A REPLY)
Explanation
The minimum amount of persistent storage that a PVC can request is 50 gigabytes. If the request is for less than 50 gigabytes, the request is rounded up to 50 gigabytes.
https://docs.cloud.oracle.com/en-us/iaas/Content/ContEng/Tasks/contengcreatingpersistentvolumeclaim.htm

NEW QUESTION: 3
You are building a container image and pushing it to the Oracle Cloud Infrastructure Registry (OCIR). You need to make sure that these 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. Set global policy of image retention to "Retain All Images".
C. In your compartment, write a policy to limit accessto the specific repository.
D. Edit the tenancy global retention policy.
Answer: D (LEAVE A REPLY)
Explanation
Deleting an Image
When you no longer need an old image or you simply want to clean up the list of image tags in a repository, you can delete images from Oracle Cloud Infrastructure Registry.
Your permissions control the images in Oracle Cloud Infrastructure Registry that you can delete. You can delete images from repositories you've created, and from repositories that the groups to which you belong have been granted accessby identity policies. If you belong to the Administrators group, you can delete images from any repository in the tenancy.
Note that as well deleting individual images , you can set up image retention policies to delete images automatically based on selection criteria you specify (see Retaining and Deleting Images Using Retention Policies).
Note:
In each region in a tenancy, there's a global image retention policy. The global image retention policy's default selection criteria retain all images so that no images are automaticallydeleted. However, you can change the global image retention policy so that images are deleted if they meet the criteria you specify. A region's global image retention policy applies to all repositories in the region, unless it is explicitly overridden byone or more custom image retention policies.
You can set up custom image retention policies to override the global image retention policy with different criteria for specific repositories in a region. Having created a custom image retention policy, you apply the custom retention policy to a repository by adding the repository to the policy. The global image retention policy no longer applies to repositories that you add to a custom retention policy.

https://docs.cloud.oracle.com/en-us/iaas/Content/Registry/Tasks/registrymanagingimageretention.htm#:~:text=In

NEW QUESTION: 4
A leading insurancefirm is hosting its customer portal in Oracle Cloud Infrastructure (OCI) Container Engine for Kubernetes with an OCI Autonomous Database. Their support team discovered a lot of SQL injection attempts and cross-site scripting attacks to the portal, which isstarting to affect the production environment.
What should they implement to mitigate this attack?
A. Network Security Lists
B. Network Security Groups
C. Network Security Firewall
D. Web Application Firewall
Answer: D (LEAVE A REPLY)
Explanation
Web Application Firewall (WAF):
Oracle Cloud Infrastructure Web Application Firewall (WAF) is a cloud-based, Payment Card Industry (PCI) compliant, global security service that protects applications from malicious and unwanted internet traffic.
WAF can protect any internet facing endpoint,providing consistent rule enforcement across a customer's applications.
WAF provides you with the ability to create and manage rules for internet threats including Cross-Site Scripting (XSS), SQL Injection and other OWASP-defined vulnerabilities. Unwantedbots can be mitigated while tactically allowed desirable bots to enter. Access rules can limit based on geography or the signature of the request.
The global Security Operations Center (SOC) will continually monitor the internet threat landscape acting asan extension of your IT infrastructure.
References:
https://docs.cloud.oracle.com/en-us/iaas/Content/WAF/Concepts/overview.htm

NEW QUESTION: 5
A pod security policy (PSP) is implemented in your Oracle Cloud Infrastructure Container Engine for Kubernetes cluster Which rule can you use to prevent a container from running as root using PSP?
A. NoPrivilege
B. RunOnlyAsUser
C. MustRunAsNonRoot
D. forbiddenRoot
Answer: C (LEAVE A REPLY)
Explanation
What is a Pod Security Policy?

A Pod Security Policy is a cluster-level resource that controlssecurity sensitive aspects of the pod specification. The PodSecurityPolicy objects define a set of conditions that a pod must run with inorder to be accepted into the system, as well as defaults for the related fields. They allow an administrator to control the following:
Privilege Escalation
These options control the allowPrivilegeEscalation container option. This bool directly controls whether the no_new_privs flag gets set on the container process. This flag will prevent setuid binaries from changing the effective user ID, and prevent files from enabling extra capabilities (e.g. it will prevent the use of the ping tool). This behavior is required to effectively enforce MustRunAsNonRoot.
example:
Require the container to run without root privileges.
rule: 'MustRunAsNonRoot'

NEW QUESTION: 6
What is the difference between blue/green and canary deployment strategies?
A. In blue/green, application Is deployed In minor increments to a select group of people. In canary, both old and new applications are simultaneously in production.
B. In blue/green, both old and new applications are in production at the same time. In canary, application is deployed Incrementally to a select group of people.
C. In blue/green, current applications are slowly replaced with new ones. In < MW y, Application Il deployed incrementally to a select group of people.
D. 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.
Answer: B (LEAVE A REPLY)
Explanation
Blue-green deployment is a technique that reduces downtime and risk by running two identical production environments called Blue and Green. At any time, only one of the environments is live, with the live environment serving all production traffic. For this example, Blue is currently live and Green is idle.
https://docs.cloudfoundry.org/devguide/deploy-apps/blue-green.html
Canary deployments are a pattern for rolling out releases to a subset of users or servers. The idea is to first deploy the change to a small subset of servers, test it, and then roll the change out to the rest of the servers. ...
Canaries were once regularly used in coal mining as an early warning system.
https://octopus.com/docs/deployment-patterns/canary-deployments

NEW QUESTION: 7
Which two "Action Type" options are NOT available in an Oracle Cloud Infrastructure (OCI) Events rule definition?
A. Notifications
B. Functions
C. Streaming
D. Email
E. Slack
Answer: D,E (LEAVE A REPLY)
Explanation
Overview ofEvents
Oracle Cloud Infrastructure Events enables you to create automation based on the state changes of resources throughout your tenancy. Use Events to allow your development teams to automatically respond when a resource changes its state.
Event Rules must also specify an action to trigger when the filter finds a matching event. Actions are responses you define for event matches. You set up select Oracle Cloud Infrastructure services that the Events service has established as actions. The resources forthese services act as destinations for matching events.
When the filter in the rule finds a match, the Events service delivers the matching event to one or more of the destinations you identified in the rule. The destination service that receives the eventthen processes the event in whatever manner you defined. This delivery provides the automation in your environment.
You can only deliver events to certain Oracle Cloud Infrastructure services with a rule. Use the following services to create actions:
Notifications
Streaming
Functions

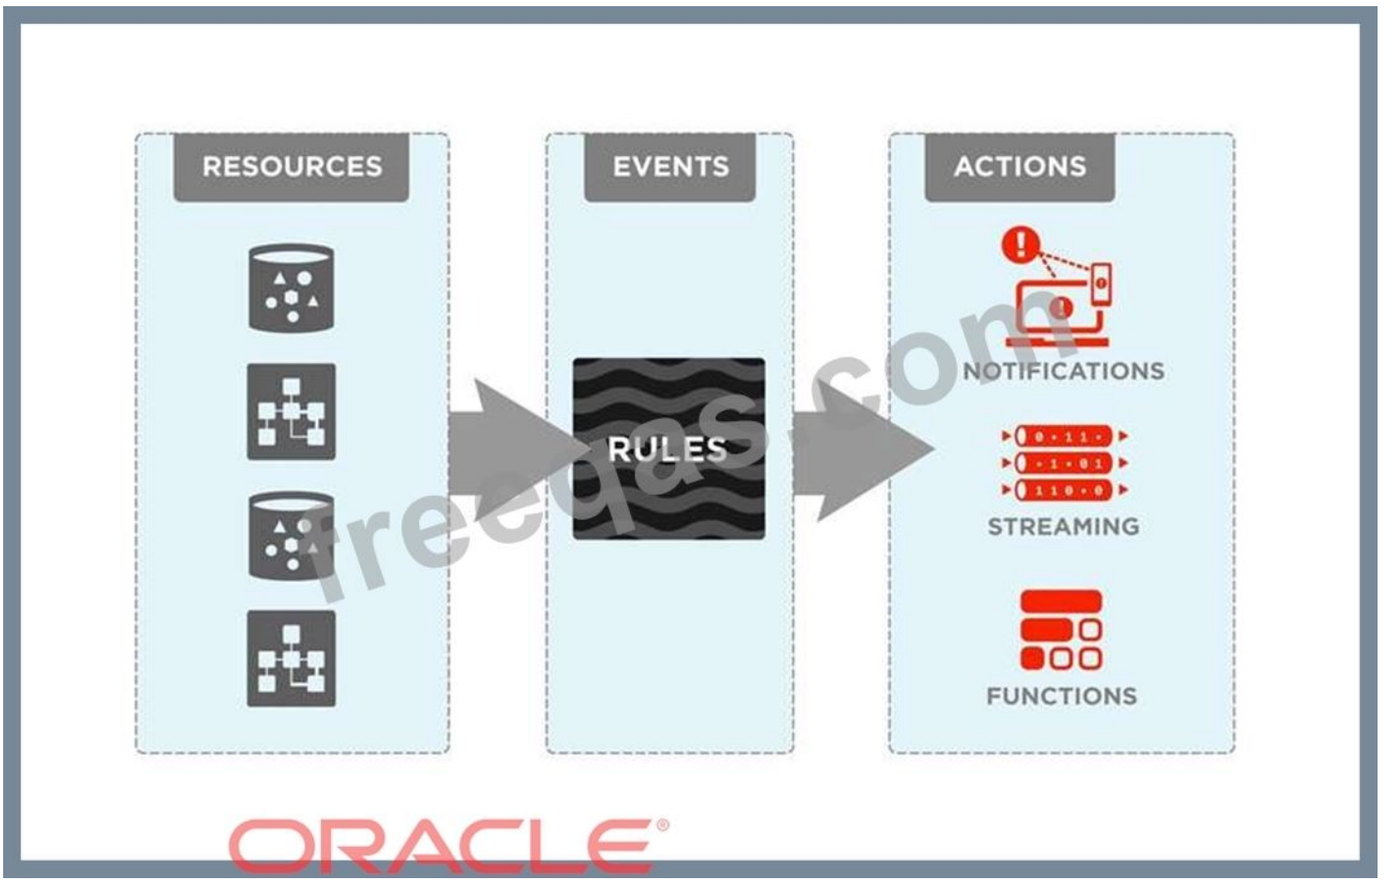

References:
https://docs.cloud.oracle.com/en-us/iaas/Content/Events/Concepts/eventsoverview.htm

NEW QUESTION: 8
With the volume of communication that can happen betweendifferent components in cloud-native applications, it is vital to not only test functionality, but also service resiliency.
Which statement is true with regards to service resiliency?
A. Resiliency is about recovering from failures without downtime or dataloss.
B. A goal of resiliency is not to bring a service to a functioning state after a failure.
C. Resiliency testing can be only done in a test environment.
D. Resiliency is about avoiding failures.
Answer: A (LEAVE A REPLY)
Explanation
Implement resilient applications:
Resiliency is the ability to (recover) from failures and continue to function. It isn't about avoiding failures but accepting the fact that failures will happen and responding to them in a way that avoids downtime or data loss. The goal of resiliency is to return the application to a fully functioning state after a failure.
References:
https://docs.microsoft.com/en-us/dotnet/architecture/microservices/implement-resilient-applications/

NEW QUESTION: 9
You are deploying an API via Oracle Cloud Infrastructure (OCI) API Gateway and you want to implement request policies to control access Which is NOT available in OCI API Gateway?
A. Limiting the number of requests sent to backend services
B. Enabling CORS (Cross-Origin Resource Sharing) support
C. Providing authentication and authorization
D. Controlling access to OCI resources
Answer: (SHOW ANSWER)
Explanation
Adding Request Policies and Response Policies to API Deployment Specifications:
You can control thebehavior of an API deployment you create on an API gateway by adding request and response policies to the API deployment specification:
a request policy describes actions to be performed on an incoming request from a caller before it is sent to a back end a response policy describes actions to be performed on a response returned from a back end before it is sent to a caller You can use request policies to:
limit the number of requests sent to back-end services
enable CORS (Cross-Origin Resource Sharing) support
provide authentication and authorization
You can add request and response policies that apply globally to all routes in an API deployment specification, and also (in some cases) request and response policies that apply only to particular routes.
Notethe following:
No response policies are currently available.
API Gateway request policies and response policies are different to IAM policies, which control access to Oracle Cloud Infrastructure resources.
You can add request and response policies to anAPI deployment specification by:
using the Console
editing a JSON file
References:
https://docs.cloud.oracle.com/en-us/iaas/Content/APIGateway/Tasks/apigatewayaddingrequestpolicies.htm

NEW QUESTION: 10
You are developing a distributed application and you need a call to a path to always return a specific JSON content deploy an Oracle Cloud Infrastructure API Gateway with the below API deployment specification.

```
{
  "routes": [{
    "path": "/hello",
    "methods": ["GET"],
    "backend": {
      "type": "_____",
      "status": 200,
      "headers": [{
        "name": "Content-Type",
        "value": "application/json"
      }],
      "body" : "{\"myjson\": \"consistent response\"}"
    }
  }]
}
```


What is the correct value for type?

A. STOCK_RESPONSE_BACKEND

B. CONSTANT_BACKEND

C. JSON_BACKEND

D. HTTP_BACKEND

Answer: A (LEAVE A REPLY)

Explanation

Adding Stock Responses as an API Gateway Back End:

You'll often want to verify that an API has been successfully deployed on an API gateway without having to set up an actual back-end service. One approach is to define a route in the API deployment specification that has a path to a 'dummy' back end. Onreceiving a request to that path, the API gateway itself acts as the back end and returns a stock response you've specified. Equally, there are some situations in a production deployment where you'll want a particular path for a route to consistentlyreturn the same stock response without sending a request to a back end. For example, when you want a call to a path to always return a specific HTTP status code in the response.

Using the API Gateway service, you can define a path to a stock response backend that always returns the same:

HTTP status code

HTTP header fields (name-value pairs)

content in the body of the response

"type": "STOCK_RESPONSE_BACKEND" indicates that the API gateway itself will act as the back end and return the stock response you define (the status code, the header fields and the body content).

References:

https://docs.cloud.oracle.com/en-us/iaas/Content/APIGateway/Tasks/apigatewayaddingstockresponses.htm

NEW QUESTION: 11

What is thecommunication method between different Cloud native applications services?

A. Complex and asynchronous

B. Basic and synchronous
C. Complex and synchronous
D. Basic and asynchronous
Answer: (SHOW ANSWER)
Explanation
What Is Cloud Native?
Cloud native technologies arecharacterized by the use of containers, microservices, serverless functions, development pipelines, infrastructure expressed as code, event-driven applications, and Application Programming Interfaces (APIs). Cloud native enables faster software developmentand the ability to build applications that are resilient, manageable, observable, and dynamically scalable to global enterprise levels.
When constructing a cloud-native application, you'll want to be sensitive to how back-end services communicate with each other. Ideally, the less inter-service communication, the better. However, avoidance isn't always possible as back-end services often rely on one another to complete an operation.
While direct HTTP calls between microservices are relatively simple to implement, care should be taken to minimize this practice. To start, these calls are always synchronous and will block the operation until a result is returned or the request times outs. What were once self-contained, independent services, able to evolve independently and deploy frequently, now become coupled to each other. As coupling among microservices increase, their architectural benefits diminish. Executing an infrequent request that makes a single direct HTTP call to another microservice might be acceptable for some systems. However, high-volume calls that invoke direct HTTP calls to multiple microservices aren't advisable. They can increase latency and negatively impact the performance, scalability, and availability of your system. Even worse, a longseries of direct HTTP communication can lead to deep and complex chains of synchronous microservices calls, shown in Figure 4-9:

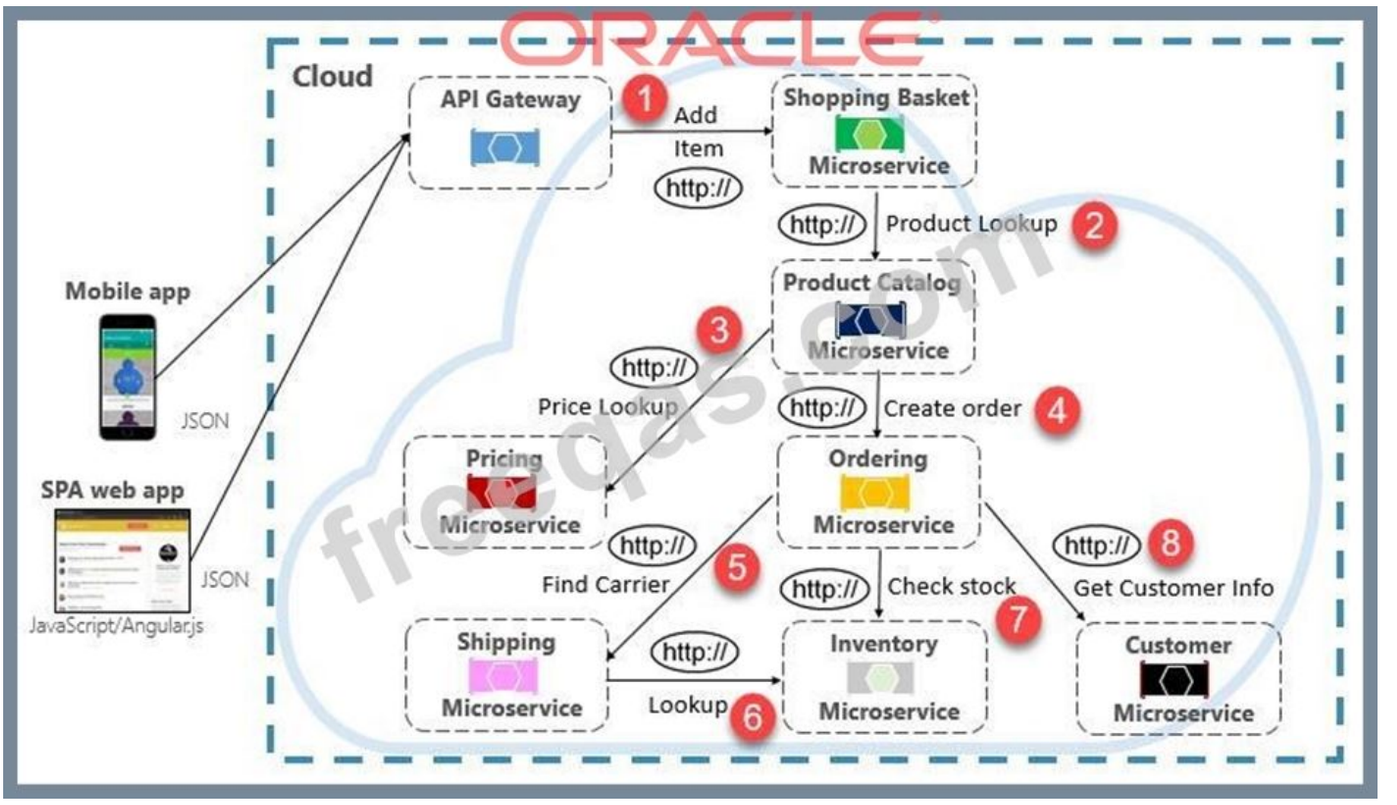

A message queue is an intermediary construct through which a producer and consumer pass a message.

Queues implement an asynchronous, point-to-point messaging pattern.

Events

Message queuing is an effective way to implement communication where a producer can asynchronously send a consumer a message.

References:

https://www.xenonstack.com/blog/cloud-native-architecture/

https://www.oracle.com/sa/cloud/cloud-native/

https://www.oracle.com/technetwork/topics/entarch/cloud-native-app-development-wp-3664668.pdf

NEW QUESTION: 12

You are working on a cloud native e-commerce application on Oracle Cloud Infrastructure (OCI). Your application architecture has multiple OCI services, including Oracle Functions. You need to trigger these functions directly from other OCI services, without having to run custom code.

Which OCI service cannot trigger your functions directly?

A. OCI Events Service

B. OCI Registry

C. OCI API Gateway

D. Oracle Integration

Answer: (SHOW ANSWER)

Explanation

Overview of Functions:

Oracle Functions is a fully managed, multi-tenant, highly scalable, on-demand, Functions-as-a-Service platform. It is built on enterprise-gradeOracle Cloud Infrastructure and powered by the Fn Project open source engine. Use Oracle Functions (sometimes abbreviated to just Functions) when you want to focus on writing code to meet business needs.
The serverless and elastic architecture of Oracle Functions means there's no infrastructure administration or software administration for you to perform. You don't provision or maintain compute instances, and operating system software patches and upgrades are applied automatically. Oracle Functions simply ensures your app is highly-available, scalable, secure, and monitored. With Oracle Functions, you can write code in Java, Python, Node, Go, and Ruby (and for advanced use cases, bring your own Dockerfile, and Graal VM).
You can invoke a function that you'vedeployed to Oracle Functions from:
- The Fn Project CLI.
- The Oracle Cloud Infrastructure SDKs.
- Signed HTTP requests to the function's invoke endpoint. Every function has an invoke endpoint.
- Other Oracle Cloud services (for example, triggered by an event in the Events service) or from external services.
so You can then deploy your code, call it directly or trigger it in response to events, and get billed only for the resources consumed during the execution.
Invoking Oracle Functions from Other OracleCloud Infrastructure Services:
You can invoke functions in Oracle Functions from other Oracle Cloud Infrastructure services. Typically, you'll want an event in another service to trigger a request to invoke a function defined in Oracle Functions.
This functionality is currently available in:
A:The Events service. For more information, see Overview of Events.
B:The Notifications service. For more information, see Notifications Overview. For a scenario, see Scenario A: Automatically Resize VMs.
C:The API Gateway service. For more information, see Adding a Function in Oracle Functions as an API Gateway Back End.
D:The Oracle Integration service, using the OCI Signature Version 1 security policy. For more information, see Configure Oracle Integration to CallOracle Cloud Infrastructure Functions with the REST Adapter in Using the REST Adapter with Oracle Integration.

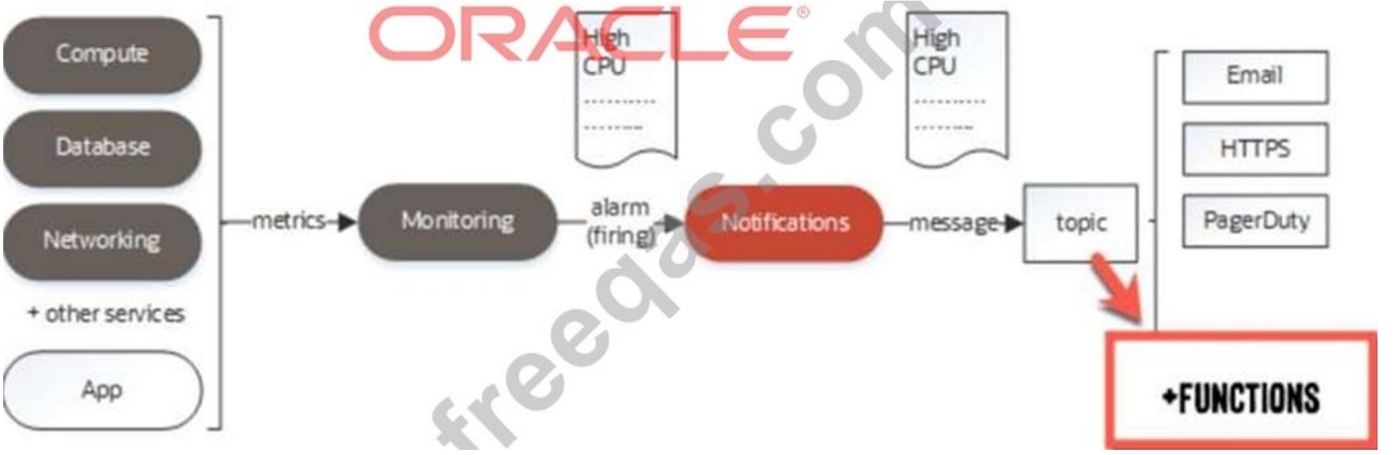

so OCI Registry services cannot trigger yourfunctions directly
References:
https://docs.cloud.oracle.com/en-us/iaas/Content/Functions/Tasks/functionsintegratingwithother.htm
https://docs.cloud.oracle.com/en-us/iaas/Content/Functions/Concepts/functionsoverview.htm
https://blogs.oracle.com/cloud-infrastructure/announcing-notifications-triggers-for-serverless-functions

NEW QUESTION: 13
Given a service deployed on Oracle Cloud infrastructure Container Engine for Kubernetes (OKE), which annotation should you add in the sample manifest file to specify a 400 Mbps load balancer?

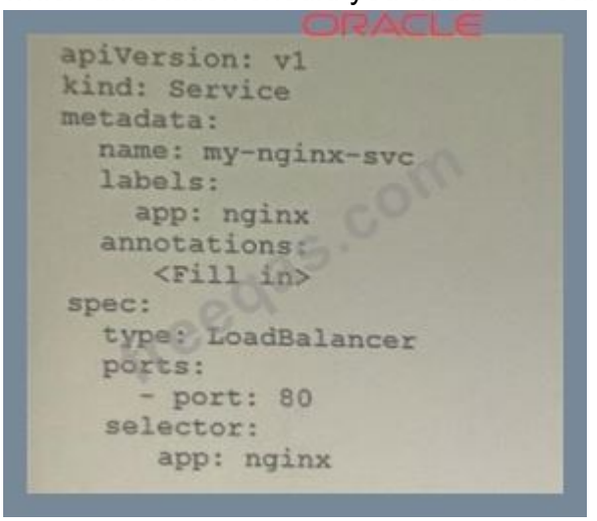

A. service.beta, kubernetes. lo/oci-load-balancer-kind: 400Mbps
B. service, beta, kubernetes. lo/oci-load-balancer-value: 4 00Mbps
C. service . beta. kubernetes . lo/oci-load-balancer-shape: 400Mbps
D. service . beta . kubernetes . lo/oci-load-balancer-size: 400Mbps
Answer: C (LEAVE A REPLY)
Explanation
The shape of an Oracle Cloud Infrastructure load balancer specifies its maximum total bandwidth (that is, ingress plus egress). By default, load balancers are created with a shape of 100Mbps. Other shapes are available, including400Mbps and 8000Mbps.
To specify an alternative shape for a load balancer, add the following annotation in the metadata section of the manifest file:
service.beta.kubernetes.io/oci-load-balancer-shape: <value>
where value is the bandwidth of the shape (for example, 100Mbps, 400Mbps, 8000Mbps).
For example:
apiVersion: v1
kind: Service
metadata:
name: my-nginx-svc
labels:

app: nginx
annotations:
service.beta.kubernetes.io/oci-load-balancer-shape: 400Mbps
spec:
type: LoadBalancer
ports:
-port: 80
selector:
app: nginx
https://github.com/oracle/oci-cloud-controller-manager/blob/master/docs/load-balancer-annotations.md

NEW QUESTION: 14
What is one of the differences between a microservice and a serverless function?
A. Microservices are used for long running operations and serverless functions for short running operations.
B. Microservices always use a data storeand serverless functions never use a data store.
C. Microservices are stateless and serverless functions are stateful.
D. Microservices are triggered by events and serverless functions are not.
Answer: A (LEAVE A REPLY)
Explanation
microservice is larger and can do more than a function. A function is a relatively small bit of code that performs only one action in response to an event.
Many microservices can run on several servers, and different instances of a specific microservice can run on several servers.
In many cases, microservices can be decomposed into a number of smaller stateless functions. The difference between microservices and functions is not simply the size. Functions are stateless, and they require no knowledge about or configuration of the underlying server-hence, the term serverless.
Microservices are best suited for long-running, complex applications that have significant resource and management requirements. You can migrate an existing monolithic application to microservices, which makes it easier to modularly develop features for the application and deploy it in the cloud. Microservices are also a good choice for building e-commerce sites, as they can retain information throughout a transaction and meet the needs of a 24/7 customer base.
On the other hand, serverless functions only execute when needed. Once the execution is over, the computing instance that runs the code decommissions itself. Serverless aligns with applications that are event driven, especially when the events are sporadic and the event processingis not resource-intensive. Serverless is a good choice when developers need to deploy fast and there are minimal application scaling concerns. For example, a good use of serverless computing is a scheduled task that needs to perform some data aggregation and will execute for just a few seconds.

References:
https://www.cloudflare.com/learning/serverless/glossary/serverless-microservice/
https://developer.oracle.com/java/fn-project-introduction.html
https://searchapparchitecture.techtarget.com/answer/When-should-I-choose-between-serverless-and-microservice

NEW QUESTION: 15
Which one of the following is NOT a valid backend-type supported by Oracle Cloud Infrastructure (OCI) API Gateway?
A. STOCK_RESPONSE_BACKEND
B. ORACLE_FUNCTIONS_BACKEND
C. ORACLE_STREAMS_BACKEND
D. HTTP_BACKEND
Answer: (SHOW ANSWER)
Explanation
In the API Gateway service, a back end is the means by which a gateway routes requests to the back-end services that implement APIs. If you add a private endpoint back end to an API gateway, you give the API gateway access to the VCN associated with that private endpoint. You can also grant an API gateway access to other Oracle Cloud Infrastructure services as back ends. For example, you could grant an API gateway access to Oracle Functions, so you can create and deploy an API that is backed by a serverless function.
API Gateway service to create an API gateway, you can create an API deployment to access HTTP and HTTPS URLs.
https://docs.cloud.oracle.com/en-us/iaas/Content/APIGateway/Tasks/apigatewayusinghttpbackend.htm API Gateway service to create an API gateway,you can create an API deployment that invokes serverless functions defined in Oracle Functions.
https://docs.cloud.oracle.com/en-us/iaas/Content/APIGateway/Tasks/apigatewayusingfunctionsbackend.htm API Gateway service, you can define a path to a stock response back end
https://docs.cloud.oracle.com/en-us/iaas/Content/APIGateway/Tasks/apigatewayaddingstockresponses.htm

NEW QUESTION: 16

In order to effectively test your cloud-native applications, you might utilize separate environments (development, testing, staging, production, etc.). Which Oracle Cloud Infrastructure (OC1) service can you use to create and manage your infrastructure?

A. OCI Compute

B. OCI Container Engine for Kubernetes

C. OCI Resource Manager

D. OCI API Gateway

Answer: C (LEAVE A REPLY)

Explanation

Resource Manager is an Oracle Cloud Infrastructure service that allows you to automate the process of provisioning your Oracle Cloud Infrastructure resources. Using Terraform, Resource Manager helps you install, configure, and manage resources through the "infrastructure-as-code" model.

References:

https://docs.cloud.oracle.com/iaas/Content/ResourceManager/Concepts/resourcemanager.htm

NEW QUESTION: 17

You are a consumer of Oracle Cloud Infrastructure (OCI) Streaming service. Which API should you use to read and process the stream?

A. ListMessages

B. GetMessages

C. GetObject

D. ReadMessages

Answer: B (LEAVE A REPLY)

Explanation

https://docs.cloud.oracle.com/en-us/iaas/Content/Streaming/Concepts/streamingoverview.htm

Building consumers to read and process messages from a stream using the GetMessages API .

How StreamingWorks:

The Streaming service provides a robust, scalable mechanism that you can use to produce and consume high volumes of data between application components.

Here's how Streaming works: a producer publishes messages to a stream, which is an append-only log. These messages are distributed among the partitions using the message's key.

Streams are divided into a number of partitions for scalability. Partitions allow you to distribute a stream by splitting messages across multiple nodes (or brokers). Each partition can be placed on a separate machine to allow multiple consumers to read a stream in parallel. Multiple consumers can read from any partition regardless of where the partition is hosted.
A consumer can read messages from one or more streams. Each message within a stream is marked with an offset value, so a consumer can pick up where it left off if it is interrupted.
You can use the Streaming service by:
A:Creating a stream using the Console or API.
B:Using a producer to publish data to the stream.
C:Building consumers to read and process stream using the GetMessages API . .
References:
https://docs.cloud.oracle.com/en-us/iaas/Content/Streaming/Concepts/streamingoverview.htm

NEW QUESTION: 18
How do you perform a rolling update in Kubernetes?
A. kubect1 rolling-update
B. kubect1 upgrade <deployment-name> -image=*image:v2
C. kubect1 update -c <container>
D. kubect1 rolling-update <deployment-name> -image=image
Answer: B (LEAVE A REPLY)
Explanation
https://docs.oracle.com/en/cloud/iaas/wercker-cloud/wercm/quickstarts/platforms/kubernetes/

NEW QUESTION: 19
As a cloud-native developer, you have written a web service for your company. You have used Oracle Cloud Infrastructure (OCI) API Gateway service to expose the HTTP backend. However, your security team has suggested that your web service should handle Distributed Denial-of-Service (DDoS) attack. You are time-constrained and you need to make sure that this is implemented as soon as possible.
What should you do in this scenario?
A. Use OCI virtual cloud network (VCN) segregation to control DDoS.
B. Use a third party service integration to implement a DDoS attack mitigation,
C. Use OCI API Gateway service and configure rate limiting.
D. Re-write your web service and implement rate limiting.
Answer: (SHOW ANSWER)
Explanation
Having created an API gateway and deployed one or more APIs on it, you'll typically want to limit the rate at which front-end clients can make requests to back-end services. For example, to:
- maintain high availability and fair use of resources by protecting back ends from being overwhelmed by too many requests
- prevent denial-of-service attacks

- constrain costs of resource consumption
- restrict usage of APIs by your customers'users in order to monetize APIs You apply a rate limit globally to all routes in an API deployment specification.
If a request is denied because the rate limit has been exceeded, the response header specifies when the request can be retried.
You can add arate-limiting request policy to an API deployment specification by:
using the Console
editing a JSON file

```
{
"requestPolicies": {
"rateLimiting": {
"rateKey": "CLIENT_IP",
"rateInRequestsPerSecond": 10
}
},
"routes":[
{
"path": "/hello",
"methods": ["GET"],
"backend": {
"type": "ORACLE_FUNCTIONS_BACKEND",
"functionId": "ocid1.fnfunc.oc1.phx.aaaaaaaaab_____xmq"
}
}
]
}
```

https://docs.cloud.oracle.com/en-us/iaas/Content/APIGateway/Tasks/apigatewaylimitingbackendaccess.htm

NEW QUESTION: 20
Which two statements accurately describe Oracle SQL Developer Web on Oracle Cloud Infrastructure (OCI) Autonomous Database?
A. It is available for databases with dedicated Exadata infrastructure only.
B. After provisioning into an OCI compute Instance, it can automatically connect to the OCI Autonomous Databases instances.
C. It is available for databases with both dedicated and shared Exadata infrastructure.
D. It provides a development environment and a data modeler interface for OCI Autonomous Databases.
E. It must be enabled via OCI Identity and Access Management policy to get access to the Autonomous Databases instances.
Answer: (SHOW ANSWER)

Explanation
Oracle SQL Developer WebOracle SQLDeveloper Web in Autonomous Data Warehouse provides a development environment and a data modeler interface for Autonomous Databases. SQL Developer Web is available for databases with both dedicated Exadata infrastructure and shared Exadata infrastructure.
https://docs.cloud.oracle.com/en-us/iaas/Content/Database/Tasks/adbtools.htm

NEW QUESTION: 21
You need to execute a script on a remote instance through Oracle Cloud Infrastructure Resource Manager.
Which option can you use?
A. Use /bin/sh with the full path to the location of thescript to execute the script.
B. It cannot be done.
C. Download the script to a local desktop and execute the script.
D. Use remote-exec
Answer: (SHOW ANSWER)
Explanation
Using Remote Exec
With Resource Manager, you can use Terraform's remote exec functionality to execute scripts or commands on a remote computer. You can also use this technique for other provisioners that require access to the remote resource.
References:
https://docs.cloud.oracle.com/en-us/iaas/Content/ResourceManager/Tasks/usingremoteexec.htm

NEW QUESTION: 22
You are developing a serverless application with Oracle Functions. Your function needs to store state in a database. Your corporate security Standards mandate encryption of secret information like database passwords.
As a function developer, which approach should you follow to satisfy this security requirement?
A. Use the Oracle Cloud Infrastructure Console and enter the password in the function configuration section in the provided input field.
B. Use Oracle Cloud Infrastructure Key Management to auto-encrypt the password. It will inject the auto-decrypted password inside your function container.
C. Encrypt the password using Oracle Cloud Infrastructure Key Management. Decrypt this password in your function code withthe generated key.
D. All function configuration variables are automatically encrypted by Oracle Functions.
Answer: C (LEAVE A REPLY)
Explanation
Oracle Functions: Using Key Management To Encrypt And Decrypt Configuration Variables Since this process involves multiple steps, I thought it would be helpful to give you an outline of the steps that we're going to take:

Create a KMS vault
Create a Master Encryption Key
Generate a Data Encryption Key
Use the DEK plaintext return value to encrypt the sensitive value (offline) Store the encrypted sensitive value as a config variable in the serverless application Store the DEK ciphertext and the initVector used to encrypt the sensitive value as Function config variables Within the function, decrypt theDEK ciphertext back into plaintext using the OCID and Cryptographic Endpoint by invoking the OCI KMS SDK Decrypt the sensitive value using the decrypted DEK plaintext and the initVector

References:
https://blogs.oracle.com/developers/oracle-functions-using-key-management-to-encrypt-and-decrypt-configuratio
https://docs.oracle.com/en/database/other-databases/essbase/19.3/essad/encrypt-values-using-kms.html

NEW QUESTION: 23
Which two statements are true for serverless computing and serverless architectures?
A. Long running tasks are perfectly suited for serverless
B. Serverless function state should never be stored externally
C. Application DevOps team is responsible for scaling
D. Serverless function execution is fully managed by a third party
E. Applications running on a FaaS (Functions as a Service) platform
Answer: (SHOW ANSWER)
Explanation
Oracle Functions is a fully managed,multi-tenant, highly scalable, on-demand, Functions-as-a-Service platform. It is built on enterprise-grade Oracle Cloud Infrastructure and powered by the Fn Project open source engine. Use Oracle Functions (sometimes abbreviated to just Functions) when youwant to focus on writing code to meet business needs.
The serverless and elastic architecture of Oracle Functions means there's no infrastructure administration or software administration for you to perform. You don't provision or maintain compute instances, and operating system software patches and upgrades are applied automatically. Oracle Functions simply ensures your app is highly-available, scalable, secure, and monitored Applications built with a serverless infrastructure will scale automatically asthe user base grows or usage increases. If a function needs to be run in multiple instances, the vendor's servers will start up, run, and end them as they are needed.
Oracle Functions is based on Fn Project. Fn Project is an open source, container native,serverless platform that can be run anywhere - any cloud or on-premises.
Serverless architectures are not built for long-running processes. This limits the kinds of applications that can cost-effectively run in a serverless architecture. Because serverlessproviders charge for the amount of time code is running, it may cost more to run an application with long-running processes in a serverless infrastructure compared to a traditional one.
https://docs.cloud.oracle.com/en-us/iaas/Content/Functions/Concepts/functionsconcepts.htm
https://www.cloudflare.com/learning/serverless/why-use-serverless/

NEW QUESTION: 24
Which concepthe following steps reference Console instructionsCloud Infrastructure Resource Manager?
A. Job
B. Stack
C. Queue
D. Plan
Answer: C (LEAVE A REPLY)
Explanation

https://docs.cloud.oracle.com/en-us/iaas/Content/ResourceManager/Concepts/resourcemanager.htm Following are brief descriptions of key concepts and the main components of Resource Manager.

CONFIGURATION

Information to codify your infrastructure. A Terraform configuration can be either a solution or a file that you write and upload.

JOB

Instructions to perform the actions defined in your configuration. Only one job at a time can run on a given stack; further, you can have only one set of Oracle Cloud Infrastructure resources on a given stack. To provision a different set of resources, you must create a separate stack and use a different configuration.

Resource Manager provides the following job types:

Plan: Parses your Terraform configuration and creates an execution plan for the associated stack. The execution plan lists the sequence of specific actions planned to provision your Oracle Cloud Infrastructure resources. The execution plan is handed off to the apply job, which then executes the instructions.

Apply. Applies the execution plan to theassociated stack to create (or modify) your Oracle Cloud Infrastructure resources. Depending on the number and type of resources specified, a given apply job can take some time. You can check status while the job runs.

Destroy. Releases resources associated with a stack. Released resources are not deleted. For example, terminates a Compute instance controlled by a stack. The stack's job history and state remain after running a destroy job. You can monitor the status and review the results of a destroy job by inspecting the stack's log files.

Import State. Sets the provided Terraform state file as the current state of the stack. Use this job to migrate local Terraform environments to Resource Manager.

STACK

The collection of Oracle Cloud Infrastructure resources corresponding to a given Terraform configuration.

Each stack resides in the compartment you specify, in a single region;however, resources on a given stack can be deployed across multiple regions. An OCID is assigned to each stack.

the following steps reference Console instructions

Create a Terraform configuration.

Create a stack.

Run a plan job, which produces an executionplan.

Review the execution plan.

If changes are needed in the execution plan, update the configuration and run a plan job again.

Run an apply job to provision resources.

Review state file and log files, as needed.

You can optionally reapply your configuration, with or without making changes, by running an apply job again.

Optionally, to release the resources running on a stack, run a destroy job.