Ultimate Guide to Prepare Free Oracle 1z0-1084-22 Exam Questions & Answer [Q30-Q49

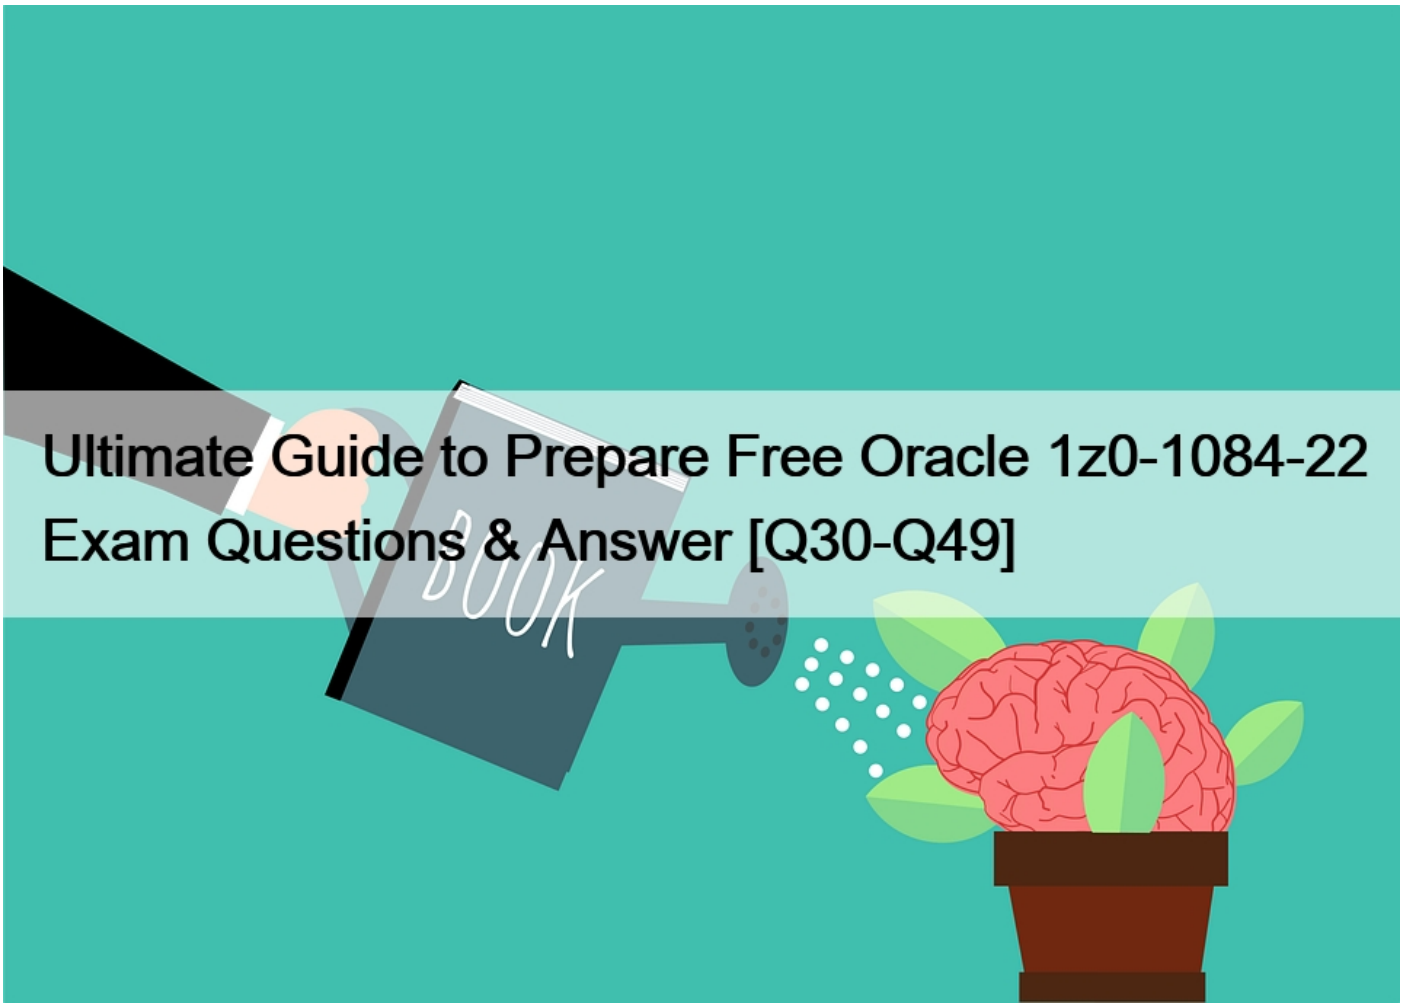

Ultimate Guide to Prepare Free Oracle 1z0-1084-22 Exam Questions and Answer
Pass Oracle 1z0-1084-22 Tests Engine pdf - All Free Dumps

NO.30 What is the minimum of storage that a persistent volume claim can obtain in Oracle Cloud Infrastructure Container Engine for Kubernetes (OKE)?
* 50 GB
* 10 GB
* 1 GB
* 1 TB
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

NO.31 A leading insurance firm 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 is starting to affect the production environment.

What should they implement to mitigate this attack?
* Network Security Lists
* Network Security Groups
* Network Security Firewall
* Web Application Firewall

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. Unwanted bots 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 as an extension of your IT infrastructure.

References:

https://docs.cloud.oracle.com/en-us/iaas/Content/WAF/Concepts/overview.htm

NO.32 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?
* kubect2 logs pod nginx
* kubect2 describe pod nginx
* kubect2 get pod nginx
* Through the Oracle Cloud Infrastructure Console

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 there been 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 or another 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/

NO.33 You are developing a serverless application with Oracle Functions. You have created a function in compartment named prod. When you try to invoke your function you get the following error.

Error invoking function. status: 502 message: dhcp options ocid1.dhcpoptions.oc1.phx.aaaaaaaac… does not exist or Oracle Functions is not authorized to use it How can you resolve this error?
* Create a policy:

Allow function-family to use virtual-network-family in compartment prod
* Create a policy:

Allow any-user to manage function-family and virtual-network-family in compartment prod
* Create a policy:

Allow service FaaS to use virtual-network-family in compartment prod
* Deleting the function and redeploying it will fix the problem
Troubleshooting Oracle Functions:

There are common issues related to Oracle Functions and how you can address them.

Invoking a function returns a FunctionInvokeSubnetNotAvailable message and a 502 error (due to a DHCP Options issue) When you invoke a function that you’ve deployed to Oracle Functions, you might see the following error message:

{“code”:”FunctionInvokeSubnetNotAvailable”,”message”:”dhcp options ocid1.dhcpoptions…….. does not exist or Oracle Functions is not authorized to use it”} Fn: Error invoking function. status: 502 message: dhcp options ocid1.dhcpoptions…….. does not exist or Oracle Functions is not authorized to use it If you see this error:

Double-check that a policy has been created to give Oracle Functions access to network resources.

Create Policies to Control Access to Network and Function-Related Resources:

Service Access to Network Resources

When Oracle Functions users create a function or application, they have to specify a VCN and a subnet in which to create them. To enable the Oracle Functions service to create the function or application in the specified VCN and subnet, you must create an identity policy to grant the Oracle Functions service access to the compartment to which the network resources belong.

To create a policy to give the Oracle Functions service access to network resources:

Log in to the Console as a tenancy administrator.

Create a new policy in the root compartment:

Open the navigation menu. Under Governance and Administration, go to Identity and click Policies.

Follow the instructions in To create a policy, and give the policy a name (for example, functions-service-network-access).

Specify a policy statement to give the Oracle Functions service access to the network resources in the compartment:

Allow service FaaS to use virtual-network-family in compartment <compartment-name> For example:

Allow service FaaS to use virtual-network-family in compartment acme-network Click Create.

Double-check that the set of DHCP Options in the VCN specified for the application still exists.

References:

https://docs.cloud.oracle.com/en-us/iaas/Content/Functions/Tasks/functionstroubleshooting.htm

https://docs.cloud.oracle.com/en-us/iaas/Content/Functions/Tasks/functionscreatingpolicies.htm

NO.34 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?

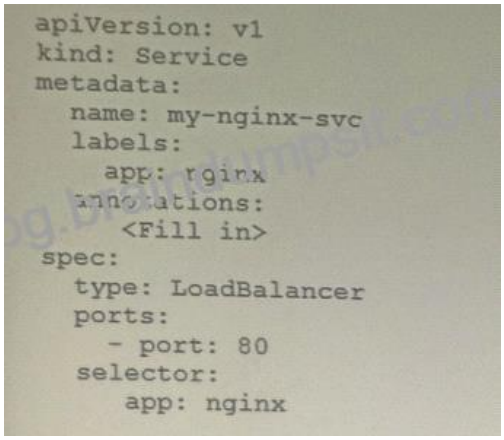

* service.beta, kubernetes. lo/oci-load-balancer-kind: 400Mbps
* service, beta, kubernetes. lo/oci-load-balancer-value: 4 00Mbps
* service . beta. kubernetes . lo/oci-load-balancer-shape: 400Mbps
* service . beta . kubernetes . lo/oci-load-balancer-size: 400Mbps

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, including 400Mbps 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

NO.35 In a Linux environment, what is the default locations of the configuration file that Oracle Cloud Infrashtructure CLI uses for profile information/
* /etc/.oci/config
* /usr/local/bin/config
* SHOME/.oci/config
* /usr/bin/oci/config
By default, the Oracle Cloud Infrastructure CLI configuration file is located at ~/.oci/config.

You might already have a configuration file as a result of installing the Oracle Cloud Infrastructure CLI.

NO.36 What is the difference between blue/green and canary deployment strategies?
* 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.
* 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.
* In blue/green, current applications are slowly replaced with new ones. In < MW y, Application ll deployed incrementally to a select group of people.
* 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.
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

NO.37 A service you are deploying to Oracle infrastructure (OCI) Container En9ine for Kubernetes (OKE) uses a docker image from a private repository Which configuration is necessary to provide access to this repository from OKE?
* Add a generic secret on the cluster containing your identity credentials. Then specify a registrycredentials property in the deployment manifest.
* Create a docker-registry secret for OCIR with API key credentials on the cluster, and specify the imagepullsecret property in the application deployment manifest.
* Create a docker-registry secret for OCIR with identity Auth Token on the cluster, and specify the image pull secret property in the application deployment manifest.
* Create a dynamic group for nodes in the cluster, and a policy that allows the dynamic group to read repositories in the same compartment.
Pulling Images from Registry during Deployment

During the deployment of an application to a Kubernetes cluster, you’ll typically want one or more images to be pulled from a Docker registry. In the application’s manifest file you specify the images to pull, the registry to pull them from, and the credentials to use when pulling the images. The manifest file is commonly also referred to as a pod spec, or as a deployment.yaml file (although other filenames are allowed).

If you want the application to pull images that reside in Oracle Cloud Infrastructure Registry, you have to perform two steps:

– You have to use kubectl to create a Docker registry secret. The secret contains the Oracle Cloud Infrastructure credentials to use when pulling the image. When creating secrets, Oracle strongly recommends you use the latest version of kubectl To create a Docker registry secret:

1- If you haven’t already done so, follow the steps to set up the cluster’s kubeconfig configuration file and (if necessary) set the KUBECONFIG environment variable to point to the file. Note that you must set up your own kubeconfig file. You cannot access a cluster using a kubeconfig file that a different user set up.

2- In a terminal window, enter:

$ kubectl create secret docker-registry <secret-name> –docker-server=<region-key>.ocir.io –docker-username='<tenancy-namespace>/<oci-username>’ –docker-password='<oci-auth-token>’ –docker-email='<email-address>’ where:

<secret-name> is a name of your choice, that you will use in the manifest file to refer to the secret . For example, ocirsecret

<region-key> is the key for the Oracle Cloud Infrastructure Registry region you’re using. For example, iad. See Availability

by Region.

ocir.io is the Oracle Cloud Infrastructure Registry name.

<tenancy-namespace> is the auto-generated Object Storage namespace string of the tenancy containing the repository from which the application is to pull the image (as shown on the Tenancy Information page). For example, the namespace of the acme-dev tenancy might be ansh81vru1zp. Note that for some older tenancies, the namespace string might be the same as the tenancy name in all lower-case letters (for example, acme-dev).

<oci-username> is the username to use when pulling the image. The username must have access to the tenancy specified by <tenancy-name>. For example, jdoe@acme.com . If your tenancy is federated with Oracle Identity Cloud Service, use the format oracleidentitycloudservice/<username>

<oci-auth-token> is the auth token of the user specified by <oci-username>. For example, k]j64r{1sJSSF-;)K8

<email-address> is an email address. An email address is required, but it doesn’t matter what you specify. For example, jdoe@acme.com

– You have to specify the image to pull from Oracle Cloud Infrastructure Registry, including the repository location and the Docker registry secret to use, in the application’s manifest file.

References:

https://docs.cloud.oracle.com/en-us/iaas/Content/Registry/Tasks/registrypullingimagesfromocir.htm

NO.38 What is the communication method between different Cloud native applications services?
* Complex and asynchronous
* Basic and synchronous
* Complex and synchronous
* Basic and asynchronous

What Is Cloud Native?

Cloud native technologies are characterized 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 development and 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 long series of direct HTTP communication can lead to deep and complex chains of synchronous microservices calls, shown in Figure 4-9:

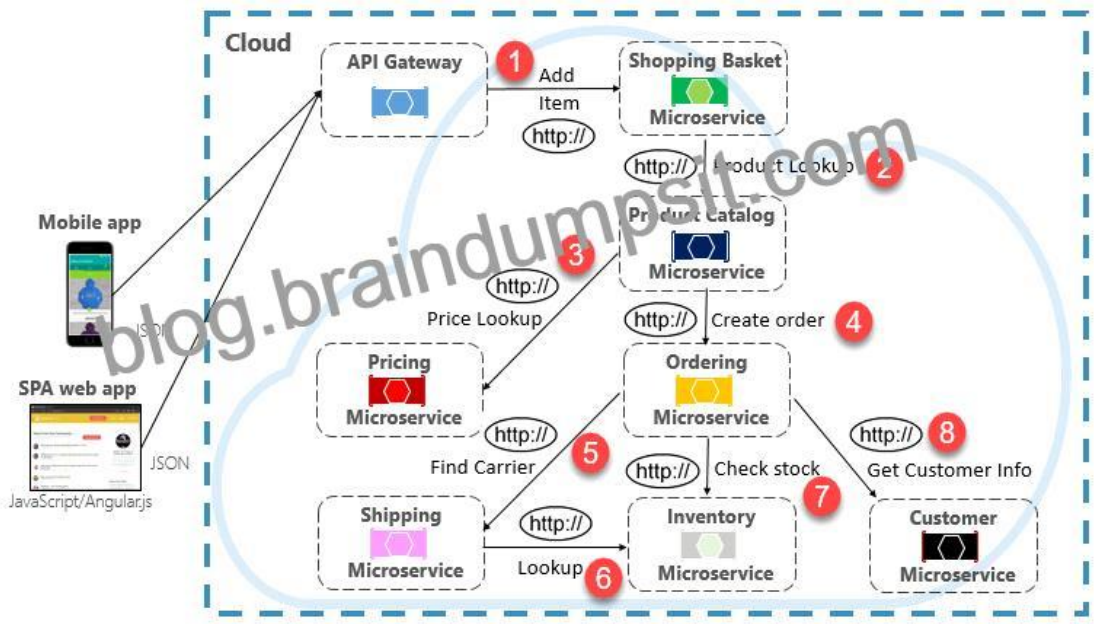

A message queue is an intermediary construct through which a producer and consumer pass a message. Queues implement an asynchronous, point-to-point messaging pattern.

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

NO.39 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?
* OCI Events Service
* OCI Registry
* OCI API Gateway
* Oracle Integration

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-grade Oracle Cloud Infrastructure and powered by the Fn Project open source engine. Use Oracle Functions (sometimes

abbreviated to just Functions) when you want to focus on writing code to meet business needs.

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’ve deployed 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 Oracle Cloud 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:

1. The Events service. For more information, see Overview of Events.

2.The Notifications service. For more information, see Notifications Overview. For a scenario, see Scenario A: Automatically Resize VMs.

3.The API Gateway service. For more information, see Adding a Function in Oracle Functions as an API Gateway Back End.

4.The Oracle Integration service, using the OCI Signature Version 1 security policy. For more information, see Configure Oracle Integration to Call Oracle Cloud Infrastructure Functions with the REST Adapter in Using the REST Adapter with Oracle Integration.

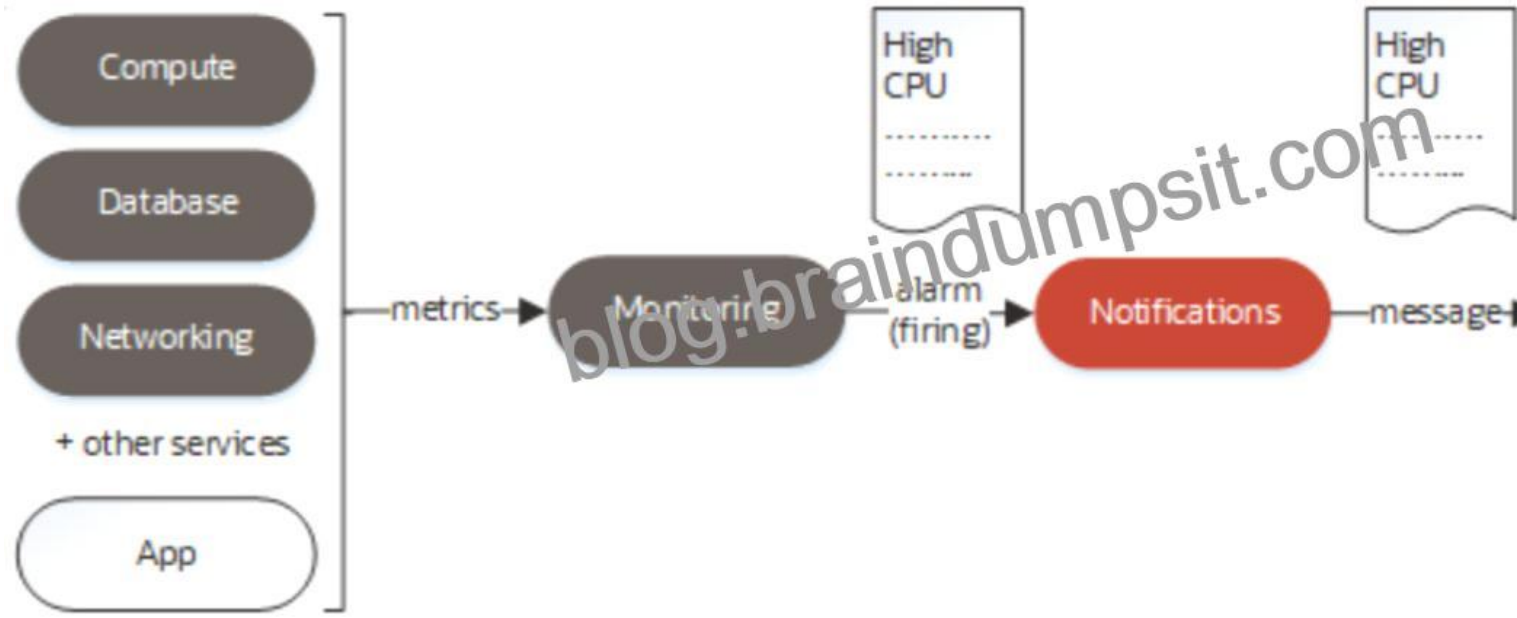

so OCI Registry services cannot trigger your functions directly

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

NO.40 A programmer Is developing a Node is application which will run in a Linux server on their on-premises data center. This application will access various Oracle Cloud Infrastructure (OC1) services using OCI SDKs.

What is the secure way to access OCI services with OCI Identity and Access Management (JAM)?
* Create a new OCI IAM user associated with a dynamic group and a policy that grants the desired permissions to OCI services. Add the on-premises Linux server in the dynamic group.
* Create an OCI IAM policy with the appropriate permissions to access the required OCI services and assign the policy to the on-premises Linux server.
* Create a new OCI IAM user, add the user to a group associated with a policy that grants the desired permissions to OCI services. In the on-premises Linux server, generate the keypair used for signing API requests and upload the public key to the IAM user.
* Create a new OCI IAM user, add the user to a group associated with a policy that grants the desired permissions to OCI services. In the on-premises Linux server, add the user name and password to a file used by Node.js authentication.
Set up an Oracle Cloud Infrastructure API Signing Key for Use with Oracle Functions:

Before using Oracle Functions, you have to set up an Oracle Cloud Infrastructure API signing key.

The instructions in this topic assume:

– you are using Linux

For more information and other options, see Required Keys and OCIDs.

The instructions below describe how to create a new ~/.oci directory, how to generate a new private key file and public key file in that ~/.oci directory, how to upload the public key to Oracle Cloud Infrastructure to create a new API signing key, and how to obtain a fingerprint for the public API key. Be aware that instructions and examples elsewhere in this documentation assume the ~/.oci directory exists and contains the private and public key files.

To set up an API signing key:

Log in to your development environment as a functions developer.

In a terminal window, confirm that the ~/.oci directory does not already exist. For example, by entering:

ls ~/.oci

Assuming the ~/.oci directory does not already exist, create it. For example, by entering:

mkdir ~/.oci

Generate a private key encrypted with a passphrase that you provide by entering:

$ openssl genrsa -out ~/.oci/<private-key-file-name>.pem -aes128 2048

where <private-key-file-name> is a name of your choice for the private key file (for example, john_api_key_private.pem).

For example:

$ openssl genrsa -out ~/.oci/john_api_key_private.pem -aes128 2048

Generating RSA private key, 2048 bit long modulus

….+++

…………………………………………………………..+++ e is 65537 (0x10001) Enter pass phrase for /Users/johndoe/.oci/john_api_key_private.pem:

References:

https://docs.cloud.oracle.com/en-us/iaas/Content/Functions/Tasks/functionssetupapikey.htm

NO.41 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?
* 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 *
* Set up the following dynamic group for your function’s OCID: Name: read-file-dg Rule: resource . id = ‘ ocid1. f nf unc. ocl -phx. aaaaaaaakeaobctakezj z5i4uj j 7g25q7sx5mvr55pms6f 4da !
* 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’
* 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’
* No policies are needed. By default, every function has read access to Object Storage buckets in the tenancy
When a function you’ve deployed to Oracle Functions is running, it can access other Oracle Cloud Infrastructure resources. For example:

– You might want a function to get a list of VCNs from the Networking service.

– You might want a function to read data from an Object Storage bucket, perform some operation on the data, and then write the modified data back to the Object Storage bucket.

To enable a function to access another Oracle Cloud Infrastructure resource, you have to include the function in a dynamic group, and then create a policy to grant the dynamic group access to that resource.

https://docs.cloud.oracle.com/en-us/iaas/Content/Functions/Tasks/functionsaccessingociresources.htm

NO.42 Your organization uses a federated identity provider to login to your Oracle Cloud Infrastructure (OCI) environment. As a developer, you are writing a script to automate some operation and want to use OCI CLI to do that. Your security team doesn’t allow storing private keys on local machines.

How can you authenticate with OCI CLI?
* Run oci setup keys and provide your credentials
* Run oci session refresh -profile <profile_name>
* Run oci session authenticate and provide your credentials
* Run oci setup oci-cli-rc -file path/to/target/file
Token-based authentication for the CLI:

Token-based authentication for the CLI allows customers to authenticate their session interactively, then use the CLI for a single session without an API signing key. This enables customers using an identity provider that is not SCIM-supported to use a federated user account with the CLI and SDKs.

Starting a Token-based CLI Session

To use token-based authentication for the CLI on a computer with a web browser:

1. In the CLI, run the following command. This will launch a web browser.

oci session authenticate

2. In the browser, enter your user credentials. This authentication information is saved to the .config file.

Validating a Token

To verify that a token is valid, run the following command:

oci session validate –config-file <path_to_config_file> –profile <profile_name> –auth security_token You should receive a message showing the expiration date for the session. If you receive an error, check your profile settings.

References:

https://docs.cloud.oracle.com/en-us/iaas/Content/API/SDKDocs/clitoken.htm

NO.43 Which two statements are true for serverless computing and serverless architectures?
* Long running tasks are perfectly suited for serverless
* Serverless function state should never be stored externally
* Application DevOps team is responsible for scaling
* Serverless function execution is fully managed by a third party
* Applications running on a FaaS (Functions as a Service) platform

Oracle Functions is a fully managed, multi-tenant, highly scalable, on-demand, Functions-as-a-Service platform. It is built on enterprise-grade Oracle Cloud Infrastructure and powered by the Fn Project open source engine. Use Oracle Functions (sometimes abbreviated to just Functions) when you want to focus on writing code to meet business needs.

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 as the 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 serverless providers 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/

NO.44 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?
* OCI Compute
* OCI Container Engine for Kubernetes
* OCI Resource Manager
* OCI API Gateway

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

NO.45 Per CAP theorem, in which scenario do you NOT need to make any trade-off between the guarantees?
* when there are no network partitions
* when the system is running in the cloud
* when the system is running on-premise
* when you are using load balancers
(1) CAP THEOREM

“CONSISTENCY, AVAILABILITY and PARTITION TOLERANCE are the features that we want in our distributed system together” Of three properties of shared-data systems (Consistency, Availability and tolerance to network Partitions) only two can be achieved at any given moment in time.

(2) In a distributed system, you can have both Consistency and Availability, except when there is a Partition:

Relaxing the consistency requirements usually makes it easier to maintain availability, but the CAP theorem is not an excuse to give up strong consistency across the board. A well-designed system can balance both availability and consistency while tolerating partitions over a range of tradeoffs, where eventual consistency is just one possibility.

References:

https://blogs.oracle.com/maa/the-cap-theorem:-consistency-and-availability-except-when-partitioned

NO.46 Which is NOT a valid option to execute a function deployed on Oracle Functions?
* Send a signed HTTP requests to the function’s invoke endpoint
* Invoke from Oracle Cloud Infrastructure CLI
* Invoke from Docker CLI
* Trigger by an event in Oracle Cloud Infrastructure Events service
* Invoke from Fn Project CLI
You can invoke a function that you’ve deployed to Oracle Functions in different ways:

Using the Fn Project CLI.

Using the Oracle Cloud Infrastructure CLI.

Using the Oracle Cloud Infrastructure SDKs.

Making a signed HTTP request to the function’s invoke endpoint. Every function has an invoke endpoint.

Each of the above invokes the function via requests to the API. Any request to the API must be authenticated by including a signature and the OCID of the compartment to which the function belongs in the request header. Such a request is referred to as a ‘signed’ request. The signature includes Oracle Cloud Infrastructure credentials in an encrypted form.

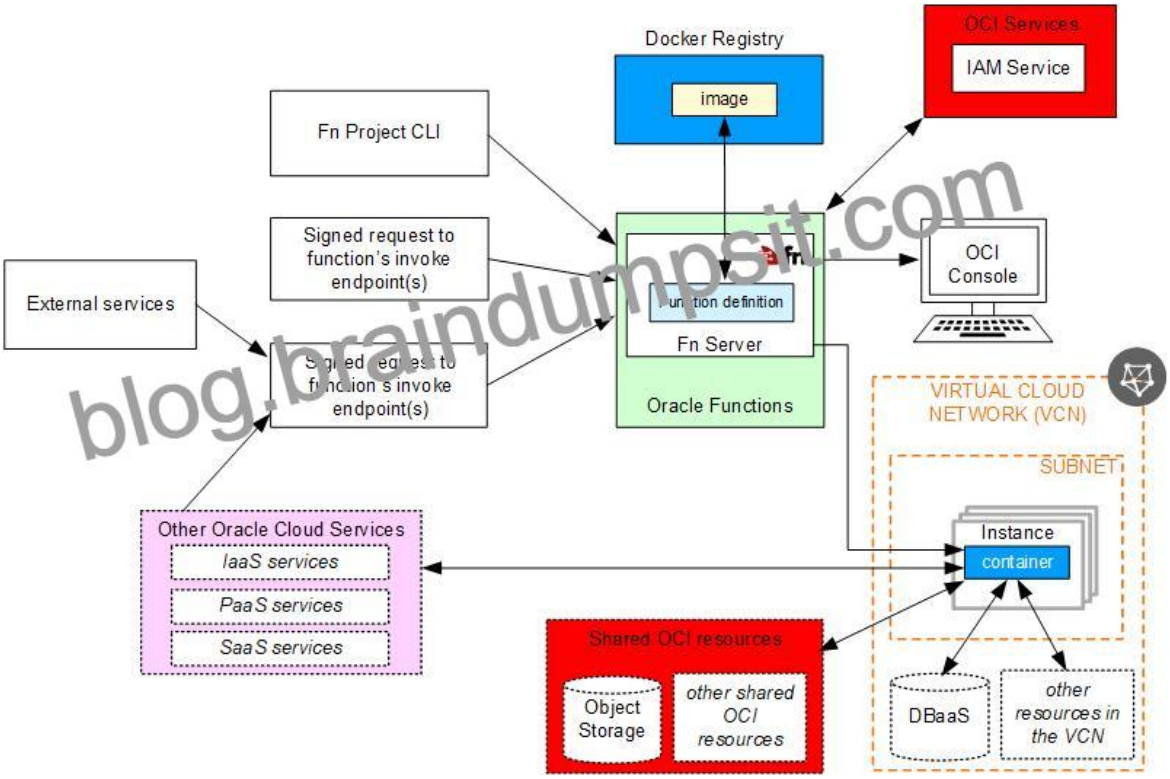

NO.47 As a cloud-native developer, you have written a web service for your company. You have used Oracle Cloud Infrastructure (OCI) API Gateway service to expose the HTTP backend. However, your security team has suggested that your web service should handle Distributed Denial-of-Service (DDoS) attack. 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?
* Use OCI virtual cloud network (VCN) segregation to control DDoS.
* Use a third party service integration to implement a DDoS attack mitigation,
* Use OCI API Gateway service and configure rate limiting.
* Re-write your web service and implement rate limiting.

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 a rate-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

NO.48 Which Oracle Cloud Infrastructure (OCI) load balancer shape is used by default in OCI container Engineer for Kubernetes?
* 400 Mbps
* 8000 Mbps
* There is no default. The shape has to be specified.
* 100 Mbps

Specifying Alternative Load Balancer Shapes

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, including 400Mbps and 8000Mbps.

SHAPE

A template that determines the load balancer’s total pre-provisioned maximum capacity (bandwidth) for ingress plus egress traffic. Available shapes include 10Mbps, 100 Mbps, 400 Mbps, and 8000 Mbps.

References:

https://docs.cloud.oracle.com/en-us/iaas/Content/ContEng/Tasks/contengcreatingloadbalancer.htm

https://docs.cloud.oracle.com/en-us/iaas/Content/Balance/Concepts/balanceoverview.htm

NO.49 You have been asked to create a stateful application deployed in Oracle Cloud Infrastructure (OCI) Container Engine for Kubernetes (OKE) that requires all of your worker nodes to mount and write data to persistent volumes.

Which two OCI storage services should you use?
* Use OCI File Services as persistent volume.
* Use GlusterFS as persistent volume.
* Use OCI Block Volume backed persistent volume.
* Use open source storage solutions on top of OCI.
* Use OCI Object Storage as persistent volume.

A PersistentVolume (PV) is a piece of storage in the cluster that has been provisioned by an administrator. PVs are volume plugins like Volumes, but have a lifecycle independent of any individual Pod that uses the PV.

A PersistentVolumeClaim (PVC) is a request for storage by a user. It is similar to a Pod. Pods consume node resources and PVCs consume PV resources.

If you intend to create Kubernetes persistent volumes, sufficient block volume quota must be available in each availability domain to meet the persistent volume claim. Persistent volume claims must request a minimum of 50 gigabytes You can define and apply a persistent volume claim to your cluster, which in turn creates a persistent volume that’s bound to the claim. A claim is a block storage volume in the underlying IaaS provider that’s durable and offers persistent storage, enabling your data to remain intact, regardless of whether the containers that the storage is connected to are terminated.

With Oracle Cloud Infrastructure as the underlying IaaS provider, you can provision persistent volume claims by attaching volumes from the Block Storage service.

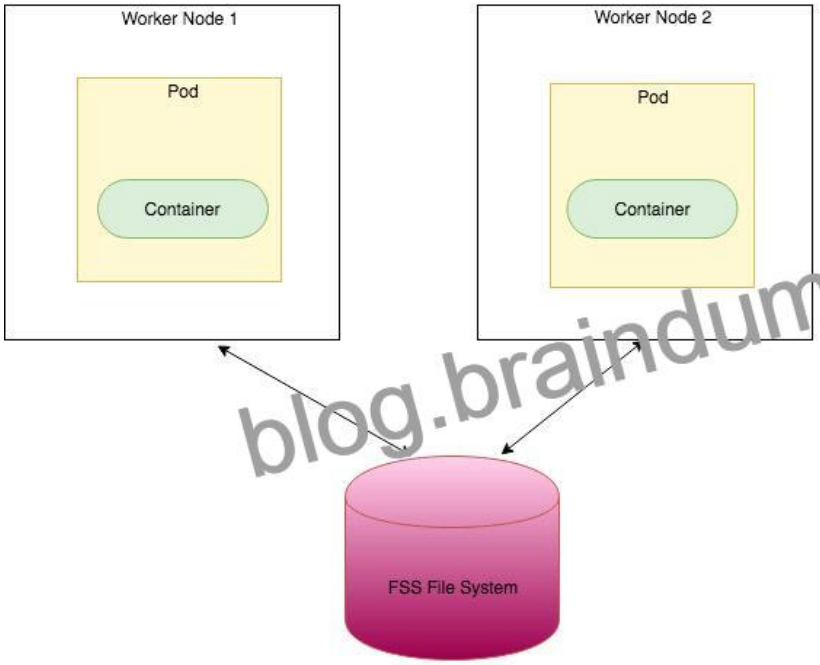

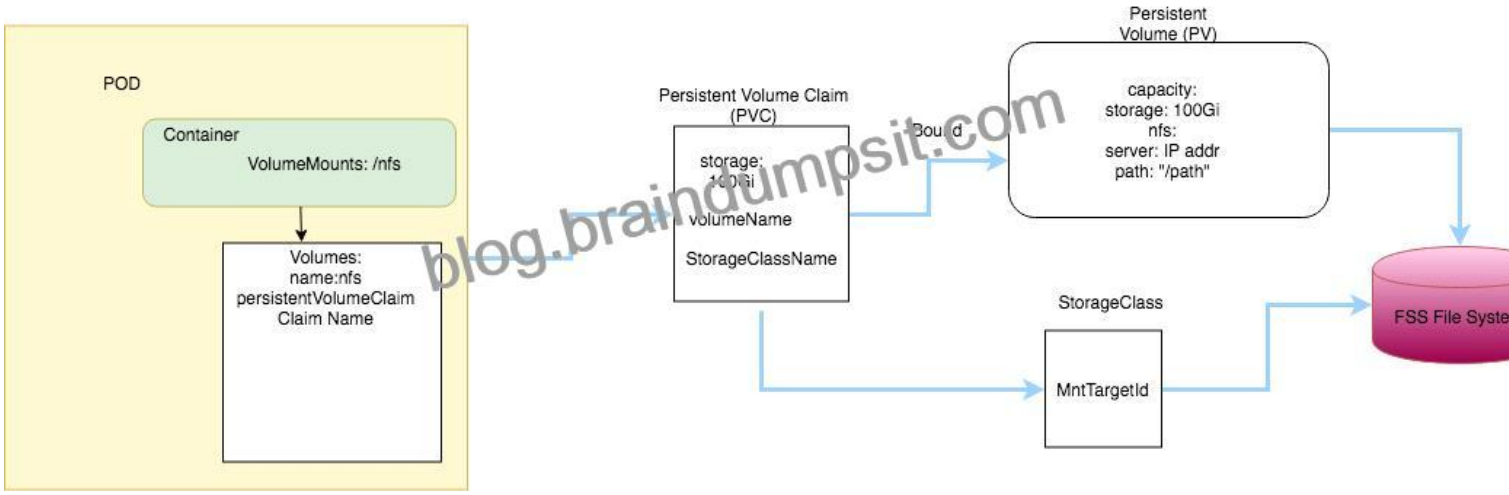

https://oracle.github.io/weblogic-kubernetes-operator/faq/oci-fss-pv/

https://kubernetes.io/docs/concepts/storage/persistent-volumes/

Online Exam Practice Tests with detailed explanations!: https://www.braindumpsit.com/1z0-1084-22_real-exam.html]