

Oracle

1Z0-1085-20 Exam

Oracle Cloud Infrastructure Foundations 2020 Associate Exam

**Questions & Answers
Demo**

Version: 8.0

Question: 1

A company has developed an eCommerce web application In Oracle Cloud Infrastructure. What should they do to ensure that the application has the highest level of resilience?

- A. Deploy the application across multiple Regions and Availability Domains.
- B. Deploy the application across multiple Availability Domains and subnet.
- C. Deploy the application across multiple Virtual Cloud Networks.
- D. Deploy the application across multiple Availability Domains and Fault Domains.

Answer: A

Explanation:

For highest level of resilience you can deploy the application between regions and distribute on availability domain and fault domains.

Reference:

<https://www.oracle.com/cloud/iaas/faq.html>

Question: 2

You have a mission-critical application which requires to be globally available at all times. Which deployment strategy should you adopt?

- A. Use multiple Fault Domains In each Availability Domain in each Region.
- B. Use multiple Availability Domains In one Region.
- C. Use multiple Fault Domains In one Region.
- D. Use multiple Fault Domains in any Availability Domain in multiple Regions.

Answer: A

Explanation:

Oracle Cloud Infrastructure is hosted in regions and availability domains. A region is a localized geographic area, and an availability domain is one or more data centers located within a region. A region is composed of one or more availability domains.

Regions are independent of other regions and can be separated by vast distances—across countries or even continents.

Availability domains are isolated from each other, fault tolerant, and very unlikely to fail simultaneously. Because availability domains do not share infrastructure such as power or cooling, or the internal availability domain network, a failure at one availability domain within a region is unlikely to impact the

availability of the others within the same region.

Fault domain is a grouping of hardware and infrastructure within an availability domain. Each availability domain contains three fault domains. Fault domains provide anti-affinity: they let you distribute your instances so that the instances are not on the same physical hardware within a single availability domain. A hardware failure or Compute hardware maintenance event that affects one fault domain does not affect instances in other fault domains. In addition, the physical hardware in a fault domain has independent and redundant power supplies, which prevents a failure in the power supply hardware within one fault domain from affecting other fault domains.

Reference:

<https://docs.cloud.oracle.com/en-us/iaas/Content/General/Concepts/regions.htm>

Question: 3

Which Oracle Cloud Infrastructure service allows you to run code without provisioning any underlying infrastructure resources?

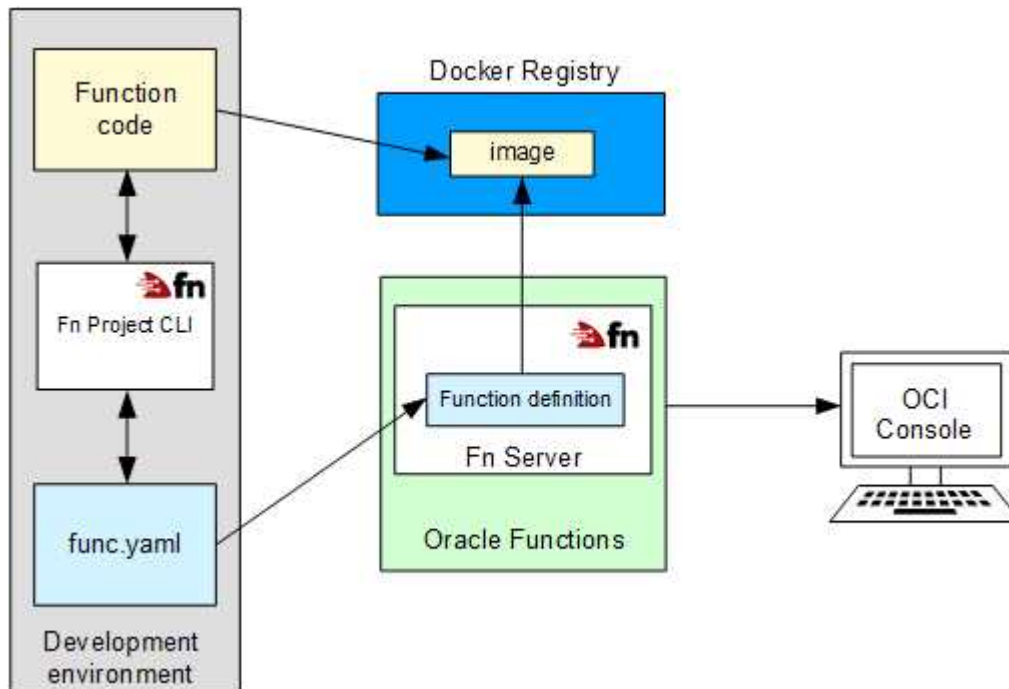
- A. Compute service
- B. Storage Gateway
- C. Oracle Container Engine for Kubernetes
- D. Oracle Functions

Answer: D

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 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 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.



Reference:

<https://docs.cloud.oracle.com/en-us/iaas/Content/Functions/Concepts/functionoverview.htm>

Question: 4

Which three methods can you use to create or modify Oracle Cloud Infrastructure (OCI) resources?

- A. REST APIs
- B. OCI desktop client
- C. Secure Shell (SSH)
- D. OCI Console
- E. Command-line Interface
- F. Remote Desktop Protocol (RDP)
- G. Serial console connection

Answer: A, D, E

Explanation:

You can create and manage resources in the following ways:

Oracle Cloud Infrastructure Console

The Console is an intuitive, graphical interface that lets you create and manage your instances, cloud networks, and storage volumes, as well as your users and permissions. See [Using the Console](#).

Oracle Cloud Infrastructure APIs

The Oracle Cloud Infrastructure APIs are typical REST APIs that use HTTPS requests and responses. See [API Requests](#).

SDKs

Several Software Development Kits are available for easy integration with the Oracle Cloud Infrastructure APIs, including SDKs for Java, Ruby, and Python. For more information, see [Developer Resources](#).

Command Line Interface (CLI)

You can use a command line interface with some services. For more information, see [Developer Resources](#).

Terraform

Oracle supports Terraform. Terraform is "infrastructure-as-code" software that allows you to define your infrastructure resources in files that you can persist, version, and share. For more information, see [Getting Started with the Terraform Provider](#).

Ansible

Oracle supports the use of Ansible for cloud infrastructure provisioning, orchestration, and configuration management. Ansible allows you to automate configuring and provisioning your cloud infrastructure, deploying and updating software assets, and orchestrating your complex operational processes. For more information, see [Getting Started with Ansible for Oracle Cloud Infrastructure](#).

Resource Manager

Resource Manager is an Oracle Cloud Infrastructure service that allows you to automate the process of provisioning your Oracle Cloud Infrastructure resources. It helps you install, configure, and manage resources using the "infrastructure-as-code" model. For more information, see [Overview of Resource Manager](#).

Reference:

https://docs.cloud.oracle.com/en-us/iaas/pdf/gsg/OCI_Getting_Started.pdf

Question: 5

Which CANNOT be used with My Oracle Support (MOS)?

- A. Add or change a tenancy administrator
- B. Request a Service Limit increase
- C. Reset the password or unlock the account for the tenancy administrator
- D. Troubleshoot your resources in an Oracle Cloud Infrastructure Free Trial account

Answer: D

Explanation:

Open a support service request with MOS option is available to paid accounts. Customers using only Always Free resources are not eligible for Oracle Support. Limited support is available to Free Tier accounts with Free Trial credits. After you use all of your credits or after your trial period ends (whichever comes first), you must upgrade to a paid account to access Oracle Support. If you choose not to upgrade and continue to use Always Free Services, you will not be eligible to raise a service request in My Oracle Support.

In addition to support for technical issues, use [My Oracle Support](#) if you need to:

- Reset the password or unlock the account for the tenancy administrator
- Add or change a tenancy administrator
- Request a [service limit](#) increase

Reference:

<https://docs.cloud.oracle.com/en-us/iaas/Content/GSG/Tasks/contactingsupport.htm>

Question: 6

Which is a key benefit of using oracle cloud infrastructure autonomous data warehouse?

- A. No username and password required
- B. Scale both CPU and Storage without downtime
- C. Apply database patches as they become available
- D. Maintain root level access to the underlying operating system

Answer: B

Explanation:

Oracle Autonomous Data Warehouse is a cloud data warehouse service that eliminates virtually all the complexities of operating a data warehouse and securing data. It automates provisioning, configuring, securing, tuning, scaling, patching, backing up, and repairing of the data warehouse. Unlike other “fully managed” cloud data warehouse solutions that only patch and update the service, it also features elastic, automated scaling, performance tuning, security, and a broad set of built-in capabilities that enable machine learning analysis, simple data loading, and data visualizations.

Data Warehouse uses continuous query optimization, table indexing, data summaries, and auto-tuning to ensure consistent high performance even as data volume and number of users grows. Autonomous scaling can temporarily increase compute and I/O by a factor of three to maintain performance. Unlike other cloud services which require downtime to scale, Autonomous Data Warehouse scales while the service continues to run.

Reference:

<https://www.oracle.com/autonomous-database/autonomous-data-warehouse/>

Question: 7

Which three components are part of Oracle Cloud Infrastructure (OCI) identity and access management service?

- A. Regional Subnets
- B. Policies
- C. Users
- D. Compute Instances
- E. Dynamic Groups
- F. Roles
- G. Virtual Cloud Networks

Answer: B, C, E

Explanation:

Components of IAM

IAM uses the components described in this section. To better understand how the components fit together, see [Example Scenario](#).

RESOURCE

The cloud objects that your company's employees create and use when interacting with Oracle Cloud Infrastructure. For example: compute instances, block storage volumes, virtual cloud networks (VCNs), subnets, route tables, etc.

USER

An individual employee or system that needs to manage or use your company's Oracle Cloud Infrastructure resources. Users might need to launch instances, manage remote disks, work with your virtual cloud network, etc. End users of your application are not typically IAM users. Users have one or more IAM credentials (see [User Credentials](#)).

GROUP

A collection of users who all need the same type of access to a particular set of resources or compartment.

DYNAMIC GROUP

A special type of group that contains resources (such as compute instances) that match rules that you define (thus the membership can change dynamically as matching resources are created or deleted). These instances act as "principal" actors and can make API calls to services according to policies that you write for the dynamic group.

NETWORK SOURCE

A group of IP addresses that are allowed to access resources in your tenancy. The IP addresses can be public IP addresses or IP addresses from a VCN within your tenancy. After you create the network source, you use policy to restrict access to only requests that originate from the IPs in the network source.

COMPARTMENT

A collection of related resources. Compartments are a fundamental component of Oracle Cloud Infrastructure for organizing and isolating your cloud resources. You use them to clearly separate resources for the purposes of measuring usage and billing, access (through the use of policies), and isolation (separating the resources for one project or business unit from another). A common approach is to create a compartment for each major part of your organization. For more information, see [Setting Up Your Tenancy](#).

TENANCY

The root compartment that contains all of your organization's Oracle Cloud Infrastructure resources. Oracle automatically creates your company's tenancy for you. Directly within the tenancy are your IAM entities (users, groups, compartments, and some policies; you can also put policies into compartments inside the tenancy). You place the other types of cloud resources (e.g., instances, virtual networks, block storage volumes, etc.) inside the compartments that you create.

POLICY

A document that specifies who can access which resources, and how. Access is granted at the group and compartment level, which means you can write a policy that gives a group a specific type of access within a specific compartment, or to the tenancy itself. If you give a group access to the tenancy, the group automatically gets the same type of access to all the compartments inside the tenancy. For more information, see [Example Scenario](#) and [How Policies Work](#). The word "policy" is used by people in different ways: to mean an individual statement written in the policy language; to mean a collection of statements in a single, named "policy" document (which has an Oracle Cloud ID (OCID) assigned to it); and to mean the overall body of policies your organization uses to control access to resources.

HOME REGION

The region where your IAM resources reside. All IAM resources are global and available across all regions, but the master set of definitions reside in a single region, the home region. You must make changes to your IAM resources in your home region. The changes will be automatically propagated to all regions. For more information, see [Managing Regions](#).

FEDERATION

A relationship that an administrator configures between an identity provider and a service provider. When you federate Oracle Cloud Infrastructure with an identity provider, you manage users and groups in the identity provider. You manage authorization in Oracle Cloud Infrastructure's IAM service. Oracle Cloud Infrastructure tenancies are federated with Oracle Identity Cloud Service by default.

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

Question: 8

A customer wants to use Oracle Cloud Infrastructure (OCI) for storing application backups which can be stored based on business needs.

Which OCI storage service can be used to meet the requirement?

- A. File Storage
- B. Block Volume
- C. Archive Storage
- D. Object Storage (standard)

Answer: D

Explanation:

Oracle Cloud Infrastructure offers two distinct storage class tiers to address the need for both performant, frequently accessed "hot" storage, and less frequently accessed "cold" storage. Storage tiers help you maximize performance where appropriate and minimize costs where possible.

1) Use Object Storage for data to which you need fast, immediate, and frequent access. Data accessibility and performance justifies a higher price to store data in the Object Storage tier.

2) Use Archive Storage for data to which you seldom or rarely access, but that must be retained and preserved for long periods of time. The cost efficiency of the Archive Storage tier offsets the long lead time required to access the data. For more information, see [Overview of Archive Storage](#).

The Oracle Cloud Infrastructure Object Storage service is an internet-scale, high-performance storage platform that offers reliable and cost-efficient data durability. The Object Storage service can store an unlimited amount of unstructured data of any content type, including analytic data and rich content, like images and videos.

With Object Storage, you can safely and securely store or retrieve data directly from the internet or from within the cloud platform. Object Storage offers multiple [management interfaces](#) that let you easily manage storage at scale. The elasticity of the platform lets you start small and scale seamlessly, without experiencing any degradation in performance or service reliability.

Object Storage is a regional service and is not tied to any specific compute instance. You can access data from anywhere inside or outside the context of the Oracle Cloud Infrastructure, as long you have internet connectivity and can access one of the [Object Storage endpoints](#). Authorization and resource limits are discussed later in this topic.

Object Storage also supports private access from Oracle Cloud Infrastructure resources in a VCN through a service gateway. A service gateway allows connectivity to the Object Storage public endpoints from private IP addresses in private subnets. For example, you can back up DB systems to an Object Storage bucket over the Oracle Cloud Infrastructure backbone instead of over the internet. You can optionally use IAM policies to control which VCNs or ranges of IP addresses can access Object Storage. See [Access to Oracle Services: Service Gateway](#) for details.

Object Storage is Always Free eligible. For more information about Always Free resources, including

additional capabilities and limitations, see [Oracle Cloud Infrastructure Free Tier](#). The following list summarizes some of the ways that you can use Object Storage.

HADOOP/BIG DATA SUPPORT

You can use Object Storage as the primary data repository for big data. Object Storage offers a scalable storage platform that lets you store large datasets and operate seamlessly on those datasets. The [HDFS Connector for Object Storage](#) provides connectivity to various big data analytic engines like Apache Spark and MapReduce. This connectivity enables the analytics engines to work directly with data stored in Object Storage. For more information, see [Hadoop Support](#).

BACKUP/ARCHIVE

You can use Object Storage to preserve backup and archive data that must be stored for an extended duration to adhere to various compliance mandates.

CONTENT REPOSITORY

You can use Object Storage as your primary content repository for data, images, logs, and video. You can reliably store and preserve this data for a long time, and serve this content directly from Object Storage. The storage scales as your data storage needs scale.

LOG DATA

You can use Object Storage to preserve application log data so that you can retroactively analyze this data to determine usage pattern and debug issues.

LARGE DATASETS

You can use Object Storage to store generated application data that needs to be preserved for future use. Pharmaceutical trials data, genome data, and Internet of Things (IoT) data are examples of generated application data that you can preserve using Object Storage.

Reference:

<https://docs.cloud.oracle.com/en-us/iaas/Content/Object/Concepts/objectstorageoverview.htm>

Question: 9

Which Oracle Cloud Infrastructure (OCI) service is best suited for running serverless apps?

- A. Oracle Functions
- B. Virtual Cloud Network
- C. Streaming
- D. Audit

Answer: A

Explanation:

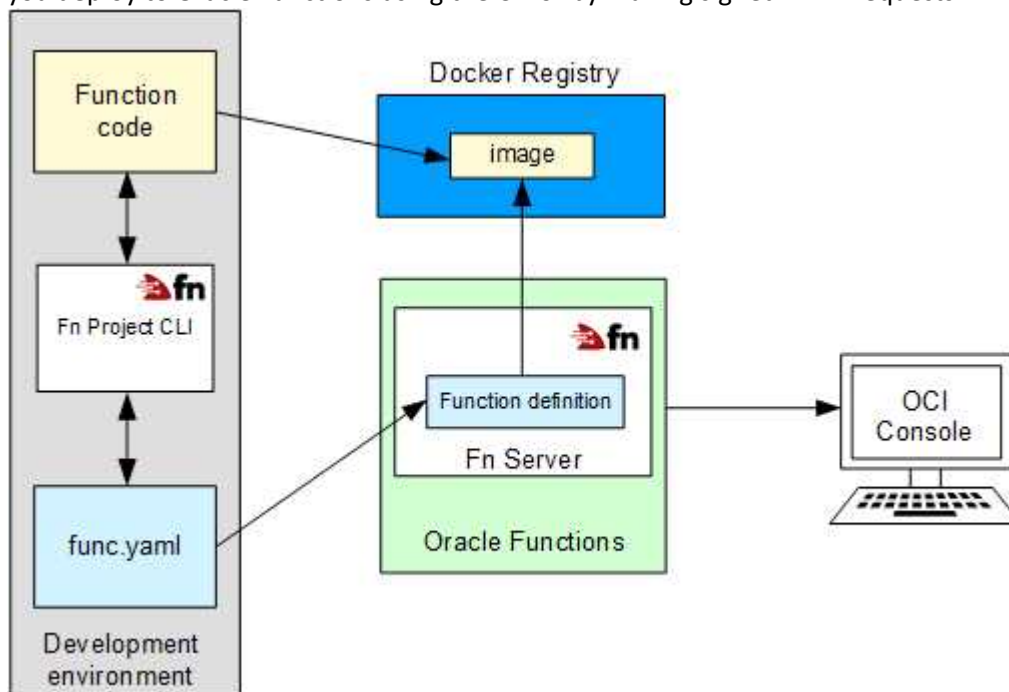
Oracle Functions is a fully managed, multi-tenant, highly scalable, on-demand, Functions-as-a-Service

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 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 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.

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. 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 Functions.

You can access Oracle Functions using the Console, a CLI, and a REST API. You can invoke the functions you deploy to Oracle Functions using the CLI or by making signed HTTP requests.



Reference:

<https://docs.cloud.oracle.com/en-us/iaas/Content/Functions/Concepts/functionoverview.htm>

Question: 10

What does compute instance horizontal scaling mean?

- A. stopping/starting the instance
- B. backing up data to object storage
- C. adding additional compute instances
- D. changing compute instance size

Answer: C

Explanation:


Cloud Horizontal Scaling refers to provisioning additional servers to meet your needs, often splitting workloads between servers to limit the number of requests any individual server is getting. In a cloud-based environment, this would mean adding additional instances instead of moving to a larger instance size.

Cloud Vertical Scaling refers to adding more CPU or memory to an existing server, or replacing one server with a more powerful server.

Reference:

<https://cloudcheckr.com/cloud-cost-management/cloud-vs-data-center-what-is-scalability-in-cloud-computing/>

Horizontal scaling means that you scale by adding more machines into your pool of resources whereas Vertical scaling means that you scale by adding more power (CPU, RAM) to an existing machine. An easy way to remember this is to think of a machine on a server rack, we add more machines across the horizontal direction and add more resources to a machine in the vertical direction.



With horizontal-scaling it is often easier to scale dynamically by adding more machines into the existing pool — Vertical-scaling is often limited to the capacity of a single machine, scaling beyond that capacity often involves downtime and comes with an upper limit.

Reference:

<https://medium.com/@abhinavkorpai/scaling-horizontally-and-vertically-for-databases-a2aef778610c>