



1Z0-1085-20^{Q&As}

Oracle Cloud Infrastructure Foundations 2020 Associate

Pass Oracle 1Z0-1085-20 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

<https://www.geekcert.com/1z0-1085-20.html>

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Oracle
Official Exam Center

-  **Instant Download** After Purchase
-  **100% Money Back** Guarantee
-  **365 Days** Free Update
-  **800,000+** Satisfied Customers





QUESTION 1

After Signing up for a new Oracle cloud Infrastructure tenancy, what would you subscribe to in order to deploy infrastructure and services in different parts of the world?

- A. Availability Domain
- B. Fault Domains
- C. Pay as you go pricing
- D. Region

Correct Answer: D

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. Most Oracle Cloud Infrastructure resources are either region-specific, such as a virtual cloud network, or availability domain-specific, such as a compute instance. Traffic between availability domains and between regions is encrypted. 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. The availability domains within the same region are connected to each other by a low latency, high bandwidth network, which makes it possible for you to provide high-availability connectivity to the internet and on-premises, and to build replicated systems in multiple availability domains for both high-availability and disaster recovery. Oracle is adding multiple cloud regions around the world to provide local access to cloud resources for our customers. To accomplish this quickly, we've chosen to launch regions in new geographies with one availability domain. As regions require expansion, we have the option to add capacity to existing availability domains, to add additional availability domains to an existing region, or to build a new region. The expansion approach in a particular scenario is based on customer requirements as well as considerations of regional demand patterns and resource availability. For any region with one availability domain, a second availability domain or region in the same country or geo-political area will be made available within a year to enable further options for disaster recovery that support customer requirements for data residency where they exist. Reference: <https://docs.cloud.oracle.com/en-us/iaas/Content/General/Concepts/regions.htm>

QUESTION 2

What is Oracle's responsibility according to the Oracle Cloud Infrastructure (OCI) shared-security model?

- A. Configuring OCI services securely
- B. Data classification and compliance
- C. Securing application workloads
- D. Security of data center facilities

Correct Answer: D

Oracle's mission is to build cloud infrastructure and platform services for your business to have effective and manageable security to run your mission-critical workloads and store your data with confidence. Oracle Cloud Infrastructure offers best-in-class security technology and operational processes to secure its enterprise cloud services. However, for you to securely run your workloads in Oracle Cloud Infrastructure, you must be aware of your security and



compliance responsibilities. By design, Oracle provides security of cloud infrastructure and operations (cloud operator access controls, infrastructure security patching, and so on), and you are responsible for securely configuring your cloud resources. Security in the cloud is a shared responsibility between you and Oracle. In a shared, multi-tenant compute environment, Oracle is responsible for the security of the underlying cloud infrastructure (such as data-center facilities, and hardware and software systems) and you are responsible for securing your workloads and configuring your services (such as compute, network, storage, and database) securely. In a fully isolated, single-tenant, bare metal server with no Oracle software on it, your responsibility increases as you bring the entire software stack (operating systems and above) on which you deploy your applications. In this environment, you are responsible for securing your workloads, and configuring your services (compute, network, storage, database) securely, and ensuring that the software components that you run on the bare metal servers are configured, deployed, and managed securely. More specifically, your and Oracle's responsibilities can be divided into the following areas:



- **Identity and Access Management (IAM):** As with all Oracle cloud services, you should protect your cloud access credentials and set up individual user accounts. You are responsible for managing and reviewing access for your own employee accounts and for all activities that occur under your tenancy. Oracle is responsible for providing effective IAM services such as identity management, authentication, authorization, and auditing.
- **Workload Security:** You are responsible for protecting and securing the operating system and application layers of your compute instances from attacks and compromises. This protection includes patching applications and operating systems, operating system configuration, and protection against malware and network attacks. Oracle is responsible for providing secure images that are hardened and have the latest patches. Also, Oracle makes it simple for you to bring the same third-party security solutions that you use today.
- **Data Classification and Compliance:** You are responsible for correctly classifying and labeling your data and meeting any compliance obligations. Also, you are responsible for auditing your solutions to ensure that they meet your compliance obligations.
- **Host Infrastructure Security:** You are responsible for securely configuring and managing your compute (virtual hosts, containers), storage (object, local storage, block volumes), and platform (database configuration) services. Oracle has a shared responsibility with you to ensure that the service is optimally configured and secured. This responsibility includes hypervisor security and the configuration of the permissions and network access controls required to ensure that hosts can communicate correctly and that devices are able to attach or mount the correct storage devices.
- **Network Security:** You are responsible for securely configuring network elements such as virtual networking, load balancing, DNS, and gateways. Oracle is responsible for providing a secure network infrastructure.
- **Client and Endpoint Protection:** Your enterprise uses various hardware and software systems, such as mobile devices and browsers, to access your cloud resources. You are responsible for securing all clients and endpoints that you allow to access Oracle Cloud Infrastructure services.



- **Client and Endpoint Protection:** Your enterprise uses various hardware and software systems, such as mobile devices and browsers, to access your cloud resources. You are responsible for securing all clients and endpoints that you allow to access Oracle Cloud Infrastructure services.
- **Physical Security:** Oracle is responsible for protecting the global infrastructure that runs all of the services offered in Oracle Cloud Infrastructure. This infrastructure consists of the hardware, software, networking, and facilities that run Oracle Cloud Infrastructure services.

Reference: https://docs.cloud.oracle.com/en-us/iaas/Content/Security/Concepts/security_overview.htm

QUESTION 3

Which gateway can be used to provide internet access to an Oracle Cloud Infrastructure compute instance in a private subnet?

- A. NAT Gateway
- B. Service Gateway
- C. Dynamic Routing Gateway
- D. Internet Gateway

Correct Answer: A

A NAT gateway gives cloud resources without public IP addresses access to the internet without exposing those resources to incoming internet connections.



Highlights

- You can add a NAT gateway to your VCN to give instances in a private subnet access to the internet.
- Instances in a private subnet don't have public IP addresses. With the NAT gateway, they can initiate connections to the internet and receive responses, but not receive inbound connections initiated from the internet.
- NAT gateways are highly available and support TCP, UDP, and ICMP ping traffic.

Overview of NAT

NAT is a networking technique commonly used to give an entire private network access to the internet without assigning each host a public IPv4 address. The hosts can initiate connections to the internet and receive responses, but not receive inbound connections initiated from the internet.

When a host in the private network initiates an internet-bound connection, the NAT device's public IP address becomes the source IP address for the outbound traffic. The response traffic from the internet therefore uses that public IP address as the destination IP address. The NAT device then routes the response to the host in the private network that initiated the connection.



Overview of NAT Gateways

The Networking service offers a reliable and highly available NAT solution for your VCN in the form of a NAT gateway.

Example scenario: Imagine you have resources that need to receive inbound traffic from the internet (for example, web servers). You also have private resources that need to be protected from inbound traffic from the internet. All of these resources need to initiate connections to the internet to request software updates from sites on the internet.

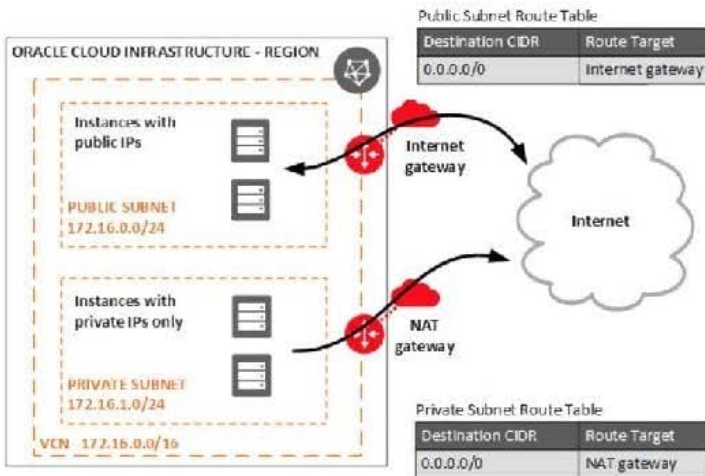
You set up a VCN and add a public subnet to hold the web servers. When launching the instances, you assign public IP addresses to them so they can receive inbound internet traffic. You also add a private subnet to hold the private instances. They cannot have public IP addresses because they are in a private subnet.

You add an internet gateway to the VCN. You also add a route rule in the public subnet's route table that directs internet-bound traffic to the internet gateway. The public subnet's instances can now initiate connections to the internet and also receive inbound connections initiated from the internet. Remember that you can use [security rules](#) to control the types of traffic that are allowed in and out of the instances at the packet level.

You add a NAT gateway to the VCN. You also add a route rule in the private subnet's route table that directs internet-bound traffic to the NAT gateway. The private subnet's instances can now initiate connections to the internet. The NAT gateway allows responses, but it does not allow connections that are initiated from the internet. Without that NAT gateway, the private instances would instead need to be in the public subnet and have public IP addresses to get their software updates.



The following diagram illustrates the basic network layout for the example. The arrows indicate whether connections can be initiated in only one direction or both.



Reference: <https://docs.cloud.oracle.com/en-us/iaas/Content/Network/Tasks/NATgateway.htm>

QUESTION 4

A new customer has logged into Oracle Cloud Infrastructure (OCI) as an administrator for the first time. The admin would like to deploy Infrastructure into a region other than their home region. What is the first Step they must take in order to accomplish this task?

- A. Use API endpoints to create resources in the desired region.
- B. Navigate to the desired region and begin creating resources.
- C. Subscribe to the desired region.
- D. File a service request for access to each additional region.

Correct Answer: C

When you sign up for Oracle Cloud Infrastructure, Oracle creates a tenancy for you in one region. This is your home region. Your home region is where your IAM resources are defined. When you subscribe to another region, your IAM resources are available in the new region, however, the master definitions reside in your home region and can only be changed there. When you subscribe your tenancy to a new region, all the policies from your home region are enforced in the new region. If you want to limit access for groups of users to specific regions, you can write policies to grant access to specific regions only. Reference: <https://docs.cloud.oracle.com/en-us/iaas/Content/Identity/Tasks/managingregions.htm>

To create an instance in another region, perform these preliminary steps:

1.
Extend your subscription to another region.
- 2.



Federate Oracle Identity Cloud Service (IDCS) from the new region with Oracle Cloud Infrastructure (OCI).

Also, when you purchase these services or sign up for a free promotion, you typically choose the data region closest to your location to access them. This becomes your primary data region. However, if required, you can extend your subscription to other geographical regions (within the same cloud account) and use the services there.

Reference:

<https://docs.oracle.com/en/cloud/paas/content-cloud/administer/create-instance-oracle-cloud-console.html>

QUESTION 5

Which should you use to distribute Incoming traffic between a set of web servers?

- A. Load Balances
- B. Internet Gateway
- C. Autoscalling
- D. Dynamic Routing Gateway

Correct Answer: A

The Oracle Cloud Infrastructure Load Balancing service provides automated traffic distribution from one entry point to multiple servers reachable from your virtual cloud network (VCN). The service offers a load balancer with your choice of a public or private IP address, and provisioned bandwidth. A load balancer improves resource utilization, facilitates scaling, and helps ensure high availability. You can configure multiple load balancing policies and application-specific health checks to ensure that the load balancer directs traffic only to healthy instances. The load balancer can reduce your maintenance window by draining traffic from an unhealthy application server before you remove it from service for maintenance. HOW LOAD BALANCING WORKS: The Load Balancing service enables you to create a public or private load balancer within your VCN. A public load balancer has a public IP address that is accessible from the internet. A private load balancer has an IP address from the hosting subnet, which is visible only within your VCN. You can configure multiple listeners for an IP address to load balance transport Layer 4 and Layer 7 (TCP and HTTP) traffic. Both public and private load balancers can route data traffic to any backend server that is reachable from the VCN. 1) Public Load Balancer To accept traffic from the internet, you create a public load balancer. The service assigns it a public IP address that serves as the entry point for incoming traffic. You can associate the public IP address with a friendly DNS name through any DNS vendor. A public load balancer is regional in scope. If your region includes multiple availability domains, a public load balancer requires either a regional subnet (recommended) or two availability domain-specific (ADspecific) subnets, each in a separate availability domain. With a regional subnet, the Load Balancing service creates a primary load balancer and a standby load balancer, each in a different availability domain, to ensure accessibility even during an availability domain outage. If you create a load balancer in two AD-specific subnets, one subnet hosts the primary load balancer and the other hosts a standby load balancer. If the primary load balancer fails, the public IP address switches to the secondary load balancer. The service treats the two load balancers as equivalent and you cannot specify which one is "primary". Whether you use regional or AD-specific subnets, each load balancer requires one private IP address from its host subnet. The Load Balancing service supplies a floating public IP address to the primary load balancer. The floating public IP address does not come from your backend subnets. If your region includes only one availability domain, the service requires just one subnet, either regional or AD-specific, to host both the primary and standby load balancers. The primary and standby load balancers each require a private IP address



from the host subnet, in addition to the assigned floating public IP address. If there is an availability domain outage, the load balancer has no failover. 2) Private Load Balancer To isolate your load balancer from the internet and simplify your security posture, you can create a private load balancer. The Load Balancing service assigns it a private IP address that serves as the entry point for incoming traffic. When you create a private load balancer, the service requires only one subnet to host both the primary and standby load balancers. The load balancer can be regional or AD-specific, depending on the scope of the host subnet. The load balancer is accessible only from within the VCN that contains the host subnet, or as further restricted by your security rules. The assigned floating private IP address is local to the host subnet. The primary and standby load balancers each require an extra private IP address from the host subnet. If there is an availability domain outage, a private load balancer created in a regional subnet within a multi-AD region provides failover capability. A private load balancer created in an AD-specific subnet, or in a regional subnet within a single availability domain region, has no failover capability in response to an availability domain outage. Reference: <https://docs.cloud.oracle.com/en-us/iaas/Content/Balance/Concepts/balanceoverview.htm>

[Latest 1Z0-1085-20 Dumps](#)

[1Z0-1085-20 PDF Dumps](#)

[1Z0-1085-20 Practice Test](#)