

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.pass2lead.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

A customer wants to deploy a customized e-commerce Web application using multiple virtual machines, block storage, databases, load balancer and web application firewall. What cloud model can be used to host this application?

- A. Software as a Service (SaaS)
- B. Platform as a Service (PaaS)
- C. Anything as a Service (XaaS)
- D. Infrastructure as a Service (IaaS)

Correct Answer: D

<https://www.oracle.com/cloud/what-is-iaas/>

What Is IaaS?

Infrastructure as a service (IaaS) is a type of cloud service model in which computing resources are hosted

in the cloud. Businesses can use the IaaS model to shift some or all of their use of on-premises or

colocated data center infrastructure to the cloud, where it is owned and managed by a cloud provider.

These infrastructure elements can include compute, network, and storage hardware as well as other

components and software. In the IaaS model, the cloud provider owns and operates the hardware and software and also owns or leases the data center. When you have an IaaS solution, you rent the resources like compute or storage, provision them when needed, and pay for the resources your organization consumes. For some resources such as compute, you'll pay for the resources you use. For others such as storage, you'll pay for capacity.

How Does IaaS Work? In a typical IaaS model, a business--which can be of any size--consumes services like compute, storage, and databases from a cloud provider. The cloud provider offers those services by hosting hardware and software in the cloud. The business will no longer need to purchase and manage its own equipment, or space to host the equipment, and the cost will shift to a pay-as-you-go model. When the business needs less, it pays for less. And when it grows, it can provision additional computing resources and other technologies in minutes. In contrast, in a traditional on-premises scenario, a business manages and maintains its own data center. The business must invest in servers, storage, software, and other technologies, and hire an IT staff or contractors to purchase, manage, and upgrade all the equipment and licenses. The data center has to be built to meet peak demand, even though sometimes workloads decline and those resources stand idle. Conversely, if the business grows quickly, the IT department might struggle to keep up. Reference: <https://www.oracle.com/in/cloud/what-is-iaas/>

QUESTION 2

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

Correct Answer: C

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-cloudcomputing/>
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/@abhinavkorpall/scaling-horizontally-and-vertically-for-databases-a2aef778610c>

QUESTION 3

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 (AD-specific) 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>

QUESTION 4

Which service is the most effective for moving large amounts of data from your on-premises to OCI?

- A. Data Transfer appliance
- B. Data Safe
- C. Internal Gateway
- D. Dynamic Routing Gateway

Correct Answer: A

QUESTION 5

You want to migrate mission-critical Oracle E- Business Suite application to Oracle Cloud Infrastructure (OCI) with full control and access to the underlying infrastructure.

Which option meets this requirement?

- A. Replace E-Business Suite with an Oracle SaaS application
- B. OCI Exadata DB Systems and OCI compute instances
- C. OCI Exadata DB Systems and Oracle Functions
- D. Oracle Exadata Cloud at customer, Storage Gateway and API Gateway

Correct Answer: B

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

[1Z0-1085-20 Study Guide](#)

[1Z0-1085-20 Exam Questions](#)