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

OCI budgets can be set on which two options?

A. Cost-tracking tags

B. Free-form tags

C. Compartments

D. Virtual Cloud Network

E. Tenancy

Correct Answer: AC

In OCI a budget can be used to set soft limits on your Oracle Cloud Infrastructure spending. You can set alerts on your budget to let you know when you might exceed your budget, and you can view all of your budgets and spending from one single place in the Oracle Cloud Infrastructure console. Budgets are set on

1.

Cost-tracking tags

2.

Compartments (including the root compartment)

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

QUESTION 2

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.

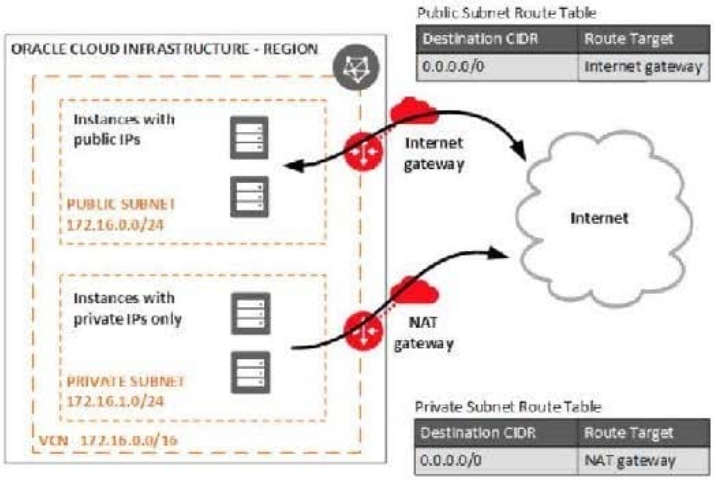

Reference: https://docs.cloud.oracle.com/en-us/iaas/Content/Network/Tasks/NATgateway.htm

QUESTION 3

How is total network throughput allocated to a Virtual Machine (VM) Instance?

A. Network bandwidth is variable

B. Network bandwidth is proportional to the number of OCPUs in the Instance shape

C. When launching a compute instance, customers may select the desired maximum network bandwidth

D. Each VM is allocated 10 Gbps of network bandwidth regardless of the selected shape

Correct Answer: B

A shape is a template that determines the number of CPUs, amount of memory, and other resources that

are allocated to an instance.

The network bandwidth is directly proportional to the number of OCPUs in the instance shape!

Flexible Shapes

A flexible shape is a shape with a customizable number of OCPUs. When you create a VM instance using the flexible shape, you select the number of OCPUs that you need for the workloads that you will run on the instance. The amount of memory, network bandwidth, and number of VNICs scale proportionately with the number of OCPUs.

The VM.Standard.E3.Flex shape, a VM standard shape, is a flexible shape.

Standard Shapes

Designed for general purpose workloads and suitable for a wide range of applications and use cases. Standard shapes provide a balance of cores, memory, and network resources. Standard shapes are available with Intel or AMD processors.

These are the bare metal standard series:

- **BM.Standard1:** X5-based standard compute. Processor: Intel Xeon E5-2699 v3. Base frequency 2.3 GHz, max turbo frequency 3.6 GHz.

  X5-based shapes availability is limited to monthly universal credit customers existing on or before November 9, 2018, in the US West (Phoenix), US East (Ashburn), and Germany Central (Frankfurt) regions.
- **BM.Standard.B1:** X6-based standard compute. Processor: Intel Xeon E5-2699 v4. Base frequency 2.2 GHz, max turbo frequency 3.6 GHz.
- **BM.Standard2:** X7-based standard compute. Processor: Intel Xeon Platinum 8167M. Base frequency 2.0 GHz, max turbo frequency 2.4 GHz.
- **BM.Standard.E2:** E2-based standard compute. Processor: AMD EPYC 7551. Base frequency 2.0 GHz, max boost frequency 3.0 GHz.
- **BM.Standard.E3:** E3-based standard compute. Processor: AMD EPYC 7742. Base frequency 2.25 GHz, max boost frequency 3.4 GHz.

VM Shapes

The following shapes are available for VMs:

- Standard Shapes
- Dense I/O Shapes
- GPU Shapes

Network bandwidth is based on expected bandwidth for traffic within a VCN.

Standard Shapes

Designed for general purpose workloads and suitable for a wide range of applications and use cases. Standard shapes provide a balance of cores, memory, and network resources. Standard shapes are available with Intel or AMD processors.

These are the VM standard series:

- **VM.Standard1:** X5-based standard compute. Processor: Intel Xeon E5-2699 v3. Base frequency 2.3 GHz, max turbo frequency 3.6 GHz.

  X5-based shapes availability is limited to monthly universal credit customers existing on or before November 9, 2018, in the US West (Phoenix), US East (Ashburn), and Germany Central (Frankfurt) regions.
- **VM.Standard.B1:** X6-based standard compute. Processor: Intel Xeon E5-2699 v4. Base frequency 2.2 GHz, max turbo frequency 3.6 GHz.
- **VM.Standard2:** X7-based standard compute. Processor: Intel Xeon Platinum 8167M. Base frequency 2.0 GHz, max turbo frequency 2.4 GHz.
- **VM.Standard.E2.1.Micro:** E2-based standard compute. Processor: AMD EPYC 7551. Base frequency

- **VM.Standard.E3:** E3-based standard compute, with a flexible number of OCPUs. Processor: AMD EPYC 7742. Base frequency 2.25 GHz, max boost frequency 3.4 GHz.

Shape	OCPU	Memory (GB)	Local Disk (TB)	Max Network Bandwidth	Max VNICs Total: Linux	Max VNICs Total: Windows
VM.Standard1.1	1	7	Block storage only	600 Mbps	2	1
VM.Standard1.2	2	14	Block storage only	1.2 Gbps	2	1
VM.Standard1.4	4	28	Block storage only	1.2 Gbps	4	1
VM.Standard1.8	8	56	Block storage only	2.4 Gbps	8	1
VM.Standard1.16	16	112	Block storage only	4.8 Gbps	16	1
VM.Standard.B1.1	1	12	Block storage only	600 Mbps	2	2
VM.Standard.B1.2	2	24	Block storage only	1.2 Gbps	2	2
VM.Standard.B1.4	4	48	Block storage only	2.4 Gbps	4	4
VM.Standard.B1.8	8	96	Block storage only	4.8 Gbps	8	8
VM.Standard.B1.16	16	192	Block storage only	9.6 Gbps	16	16
VM.Standard2.1	1	15	Block storage only	1 Gbps	2	2

Reference: https://docs.cloud.oracle.com/en-us/iaas/Content/Compute/References/computeshapes.htm

QUESTION 4

You are required to host several files in a location that can be publicly accessible from anywhere in the world. Which Oracle Cloud Infrastructure (OCI) service should you use?

A. OCI Object Storage

B. Oracle Functions

C. OCI Block Volume

D. OCI File Storage

E. OCI Storage Gateway

Correct Answer: A

QUESTION 5

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

A. Load Balances

B. Internet Gateway

C. Autoscallng

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 VCE Dumps