Verified & Latest 1z0-1085-23 Dump Q&As with Correct Answers [Q82-Q97

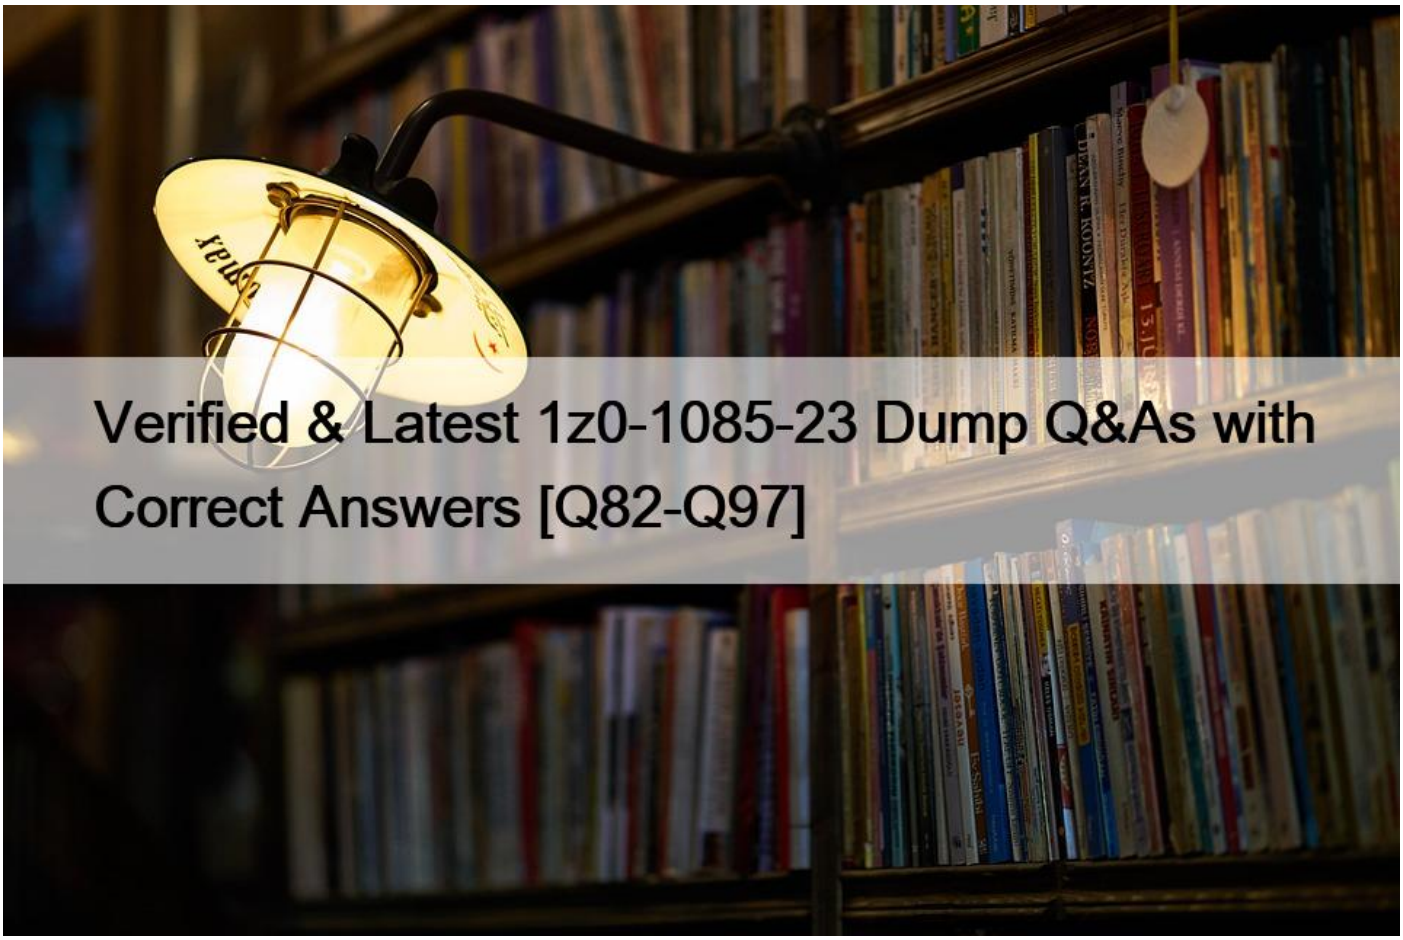

Verified & Latest 1z0-1085-23 Dump Q&As with Correct Answers
Latest 1z0-1085-23 dumps - Instant Download PDF

QUESTION 82

Which statement about the Oracle Cloud Infrastructure (OCI) shared-security model is true?
* You are responsible for securing all data that you place in OCI
* You are not responsible for any aspect of security in OCI
* You are responsible for securing the hypervisor within OCI compute service
* You are responsible for managing security controls within the physical OCI network
Explanation

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.

The responsibilities can be divided as:

QUESTION 83

Which is NOT considered a security resource within Oracle Cloud Infrastructure?
* Network Security Group
* Web Application Firewall
* File Storage Service
* Security Lists

Oracle Cloud Infrastructure File Storage service provides a durable, scalable, secure, enterprise-grade network file system. You can connect to a File Storage service file system from any bare metal, virtual machine, or container instance in your Virtual Cloud Network (VCN).

You can control the access of the file system from FSS by applying some security rules and others but the services it self not related to security but it related to shared storage Reference:

https://docs.cloud.oracle.com/en-us/iaas/Content/File/Concepts/filestorageoverview.htm

QUESTION 84

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?
* Replace E-Business Suite with an Oracle SaaS application
* OCI Exadata DB Systems and OCI compute instances
* OCI Exadata DB Systems and Oracle Functions
* Oracle Exadata Cloud at customer, Storage Gateway and API Gateway

QUESTION 85

Which component of the Oracle Cloud Infrastructure Networking service allows resources in a VCN to access Oracle Cloud Services without traversing the public internet?
* Network Address translation (NAT) Gateway
* Service gateway
* Internet gateway
* Dynamic Routing gateway (DRG)

A service gateway is the component of the Oracle Cloud Infrastructure Networking service that allows resources in a VCN to access Oracle Cloud Services without traversing the public internet. A service gateway provides a secure and private connection between a VCN and supported Oracle services, such as Object Storage, Autonomous Database, or Functions.

QUESTION 86

Which TWO are valid targets for setting up Oracle Cloud Infrastructure (OCI) budgets?

(Choose all correct answers)
* Budget tag
* Tenancy
* Identity and Access Management (IAM) group
* Compartment
* Cost-tracking tag

Compartment and Cost-tracking tag are two valid targets for setting up OCI budgets. Budgets are set on cost-tracking tags or compartments (including the root compartment) to track all spending in that cost-tracking tag or for that compartment and its children2. You can create alerts on your budget to receive email notifications based on actual or forecasted spending thresholds2.

QUESTION 87

What is the frequency of OCI usage report generation?
* Weekly
* Monthly
* Annually
* Daily

A usage report is a comma-separated value (CSV) file that can be used to get a detailed breakdown of resources in Oracle Cloud Infrastructure for audit or invoice reconciliation.

The usage report is automatically generated daily, and is stored in an Oracle-owned Object Storage bucket. It contains one row per each Oracle Cloud Infrastructure resource (such as instance, Object Storage bucket, VNIC) per hour along with consumption information, metadata, and tags. Usage reports generally contain 24 hours of usage data, although occasionally a usage report may contain late-arriving data that is older than 24 hours.

Usage reports are retained for one year.

Reference:

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

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

QUESTION 88

Which of these Is defined as a qualifier used for filtering group metric data presented in the Oracle Cloud Infrastructure (OCI) Monitoring service?
* Indicators
* Namespaces
* Dimensions
* Facts

Dimensions are defined as qualifiers used for filtering group metric data presented in the OCI Monitoring service. Dimensions are key-value pairs that provide additional information about a metric, such as resource ID, region, availability domain, fault domain, shape, image, etc. Dimensions allow you to narrow down the scope of your metric queries and alarms by specifying which resources or attributes you want to monitor. For example, you can use dimensions to filter the CPU utilization metric by instance ID or shape.

QUESTION 89

Which is NOT available to you whenever Oracle Cloud Infrastructure creates or resolves an incident?
* Twitter notifications
* Text Message notifications
* Email notifications
* Webhook notifications

The Oracle Cloud Infrastructure Notifications service broadcasts messages to distributed components through a publish-subscribe pattern, delivering secure, highly reliable, low latency and durable messages for applications hosted on Oracle Cloud Infrastructure and externally. Use Notifications to get notified when event rules are triggered or alarms are breached, or to directly publish a message.

Messages sent out as email by the Oracle Cloud Infrastructure Notifications service are processed and delivered through Oracle resources Reference:

https://docs.cloud.oracle.com/en-us/iaas/Content/Notification/Concepts/notificationoverview.htm

QUESTION 90

Which feature is NOT a component of Oracle Cloud Infrastructure (OCI) Identity and Access management service?
* User Credentials
* Network Security Group
* Federation
* Policies

QUESTION 91

Which is NOT a supported workload type for Oracle Autonomous Database?
* Transaction Processing
* APEX
* Data Warehouse
* JSON
* MySQL

QUESTION 92

Which should you use to distribute Incoming traffic between a set of web servers?
* Load Balances
* Internet Gateway
* Autoscallng
* Dynamic Routing Gateway

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.

QUESTION 93

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?
* Software as a Service (SaaS)
* Platform as a Service (PaaS)
* Anything as a Service (XaaS)
* Infrastructure as a Service (IaaS)
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.

QUESTION 94

Which THREE are capabilities of the Oracle Cloud Infrastructure (OCI) Data Catalog service? (Choose all correct answers)
* It provides a repository of searchable metadata.
* It enables enrichment of the metadata.
* It is an alternative to Autonomous Data warehouse.
* It has an accelerated library to quickly build analytics models.
* It runs Spark Jobs at scale.
* It can automate harvesting of data.
The following are capabilities of the OCI Data Catalog service:

It provides a repository of searchable metadata. Data Catalog harvests metadata from data sources across the OCI ecosystem and on-premises to create an inventory of data assets. This helps data consumers easily find the data they need for analytics4.

It enables enrichment of the metadata. Data Catalog allows users to add business metadata such as business terms, tags, custom properties, and annotations to data assets. This helps provide more insight and context into the data5.

It can automate harvesting of data. Data Catalog supports on-demand or schedule-based automatic harvesting to ensure the data catalog always has up-to-date information4.

QUESTION 95

Which statement is NOT valid regarding the Oracle Cloud Infrastructure (OCI) Block Volume service?
* You can expand an existing block volume in place with online resizing.
* You can decrease the size of a blockvolume.
* You can clone an existing block volume to a new, larger volume.
* You can increase the size of a block volume.

This statement is not valid regarding the OCI Block Volume service. You cannot decrease the size of a block volume once it is created. You can only increase the size of a block volume, either by creating a new, larger volume from a backup or a clone of the original volume, or by expanding an existing block volume in place with online resizing12 The other statements are valid regarding the OCI Block Volume service. You can expand an existing block volume in place with online resizing, which allows you to increase the size and performance of a block volume without detaching it from an instance or interrupting I/O operations. You can also clone an existing block volume to a new, larger volume, which creates an exact point-in-time copy of the source volume and preserves all the data and properties of the source volume, except for the size and performance. You can also increase the size of a block volume by creating a new, larger volume from a backup of the original volume, which restores all the data from the backup to a new volume with a different size and performance12

QUESTION 96

Which TWO are valid regarding Oracle Cloud Infrastructure (OCI) Virtual Cloud Network (VCN) peering?

(Choose all correct answers)
* Peered VCNs can have overlapping classless inter-domain routing (CIDR).
* Peered VCNs can exist in the same OCI region.
* Peered VCNs can exist in different OCI regions.
* A VCN peering connection is a VPN based connection.
* Peered VCNs need to be part of the same OCI tenancy.

VCN peering is the process of connecting multiple virtual cloud networks (VCNs) so that their resources can communicate using private IP addresses. There are four types of VCN peering: 1 Local VCN peering: This is the process of connecting two VCNs in the same region and tenancy. This type of peering uses local peering gateways (LPGs) in each VCN to establish a logical connection.

Remote VCN peering: This is the process of connecting two VCNs in different regions, either in the same tenancy or different tenancies. This type of peering uses remote peering gateways (RPGs) in each VCN to establish a logical connection.

Peering VCNs in the same region through a dynamic routing gateway (DRG): This is the process of connecting multiple VCNs in the same region and tenancy by attaching them to a common DRG. A DRG is a virtual router that provides a path for private network traffic between your VCN and other networks.

Peering VCNs in different regions through a DRG: This is the process of connecting multiple VCNs in different regions and tenancies by attaching them to a common DRG and using IPSec VPN or FastConnect to connect the DRGs across regions.

Peered VCNs cannot have overlapping CIDRs, as this would cause routing conflicts and ambiguity. Peered VCNs need to have unique CIDRs that do not overlap with each other or with any other network that they need to communicate with1 A VCN peering connection is not a VPN-based connection. A VPN-based connection is a secure and encrypted connection between your on-premises network and your OCI VCN over the public internet by using IPSec VPN or FastConnect. A VPN-based connection requires an internet gateway or a DRG in your VCN and a customer-premises equipment (CPE) device in your on-premises network2 Peered VCNs do not need to be part of the same OCI tenancy, as long as they are in different regions. Remote VCN

peering supports cross-tenancy connections, meaning that you can peer a VCN in one tenancy with a VCN in another tenancy, as long as they are subscribed to the same regions and have proper IAM policies to allow peering. Local VCN peering and peering through a DRG only support intra-tenancy connections, meaning that you can only peer VCNs within the same tenancy1

QUESTION 97

Which Oracle Cloud Infrastructure (OCI) service is best suited for running serverless apps?
* Oracle Functions
* Virtual Cloud Network
* Streaming
* Audit

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.

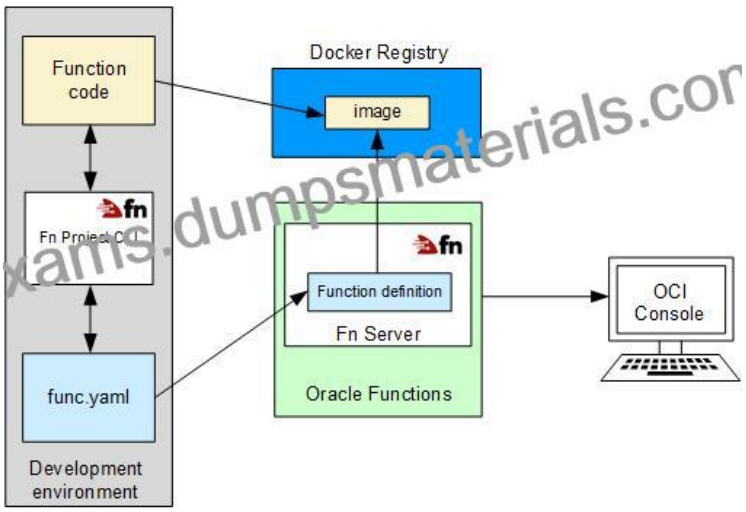

The Ultimate Oracle 1z0-1085-23 Dumps PDF Review: https://www.dumpsmaterials.com/1z0-1085-23-real-torrent.html]