

DUMPS ARENA

Oracle Cloud Data Management 2023 Foundations Associate

Oracle 1z0-1105-23

Version Demo

Total Demo Questions: 9

Total Premium Questions: 90

Buy Premium PDF

<https://dumpsarena.co>

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

Topic Break Down

Topic	No. of Questions
Topic 1, Data Management Fundamentals	24
Topic 2, Oracle Cloud Infrastructure Data Management Services	13
Topic 3, Oracle Cloud Infrastructure Database Services	50
Topic 4, Oracle Cloud Infrastructure Analytics Services	2
Topic 5, Mix Questions	1
Total	90

QUESTION NO: 1

Which Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) capability helps you to organize multiple users into teams?

- A. Users
- B. Roles
- C. Policies
- D. Groups

ANSWER: D

Explanation:

Groups is the OCI IAM capability used to organize multiple users into teams. In OCI, a group is a collection of users who require the same type of access to a set of cloud resources. Rather than managing access individually for every user, an administrator adds users to a group and writes IAM policies that grant the necessary permissions to that group. Members inherit the permissions allowed by the policies applicable to their group, making access administration more consistent and scalable as teams grow or change. For example, users on an application-development team can be placed in one group and granted the permissions needed to work with resources in a specific compartment. OCI IAM groups therefore provide the central mechanism for representing teams and applying common authorization requirements across their members. Oracle's IAM documentation defines groups as collections of users that need the same type of access to particular resources and recommends using policies to assign access to groups. See the [Oracle Cloud Infrastructure documentation on managing groups](#) and the [OCI IAM overview](#).

QUESTION NO: 2

What connections to multiple public clouds are provided through 3rd party connectivity?

- A. Megaport Cloud Router (MCR) and Virtual Cross Connect (VXC)
- B. Layer 3 and FastConnect
- C. Virtual Cloud Networks (VCNs) and Amazon Virtual Private Cloud (VPC)
- D. FastConnect and Direct Connect

ANSWER: A

Explanation:

Megaport Cloud Router (MCR) and Virtual Cross Connect (VXC) are the third-party connectivity components used to establish multicloud connectivity. Megaport Cloud Router is a virtual Layer 3 routing service deployed on Megaport's software-defined network. It allows an organization to route traffic between cloud environments without requiring physical routing equipment in a colocation facility. A Virtual Cross Connect is the logical, private connection provisioned across the Megaport network between the router and a cloud on-ramp or service endpoint. Together, these services enable a customer to connect OCI with other supported public-cloud environments through a common third-party fabric, while retaining private connectivity rather than sending intercloud traffic over the public internet. Oracle documents Megaport as a FastConnect partner and explains that partner connectivity can provide private access to OCI through the partner's network. Megaport's OCI guidance further describes using MCR and VXCs to attach OCI and other cloud services to the same virtual router. See [Oracle FastConnect Partners documentation](#) and [Megaport Cloud Router for Oracle Cloud](#).

QUESTION NO: 3

Oracle Lakehouse provides extensive data processing capabilities, to accommodate the migration of existing architectures or the creation of new, agile ones. Which service does NOT provide data processing on Oracle Cloud Infrastructure (OCI) Lakehouse?

- A. NoSQL Service
- B. Autonomous Database
- C. Big Data Service
- D. OCI Data Flow

ANSWER: A

Explanation:

NoSQL Service is correct because it is not identified as an OCI Lakehouse data-processing service. Oracle positions the Lakehouse architecture as an open, integrated platform that supports processing and analysis of data held in a lakehouse. Its processing choices include Autonomous Database for SQL-based processing, Big Data Service for Hadoop and Spark workloads, and OCI Data Flow for serverless Apache Spark workloads. These services are intended to execute transformations, queries, analytics, and large-scale processing workloads that operate across lakehouse data sources.

In contrast, Oracle NoSQL Database Cloud Service is a fully managed NoSQL database designed for low-latency, predictable application data access using flexible schemas and key-value/document-oriented data models. It is an operational database service rather than a Lakehouse processing engine. Although applications can use NoSQL data alongside broader cloud architectures, the service itself is not one of the OCI Lakehouse processing capabilities described by Oracle. Oracle's Lakehouse overview is available at [Oracle Lakehouse](#), and Oracle's service overview for the database is available at [Oracle NoSQL Database Cloud Service](#).

QUESTION NO: 4

What is NOT a workload that can be selected when provisioning an Autonomous Database?

- A. Data Warehouse
- B. MySQL
- C. APEX
- D. JSON

ANSWER: B

Explanation:

MySQL is correct because it is not one of the workload types available when creating an Oracle Autonomous Database. Autonomous Database is Oracle Database-based and lets the administrator choose a workload configuration designed for a particular Oracle Database use case. The available choices include Data Warehouse for analytical workloads, APEX for low-code application development, and JSON for document-centric applications using Oracle Database's JSON capabilities. Those selections configure the Autonomous Database environment with workload-appropriate service settings and features.

MySQL, by contrast, is a separate database engine and Oracle Cloud service offering. Oracle provides managed MySQL through MySQL Database Service, including HeatWave capabilities, rather than as an Autonomous Database workload selection. Therefore, a user provisioning an Autonomous Database cannot choose MySQL from the workload list; they would provision the relevant MySQL service instead. Oracle's Autonomous Database documentation describes the available workload types and creation choices, while the MySQL HeatWave documentation describes MySQL Database Service as its own OCI service. See [Oracle Autonomous Database documentation](#) and [Oracle MySQL HeatWave documentation](#).

QUESTION NO: 5

What are two characteristics of CI/CD?

- A. It prevents developers from merging their code changes back to the shared branch for security.
- B. It introduces automation into all stages of app/database development.

C. Development changes are automatically tested on pushes, commits, and merges.

D. It requires Liquibase

ANSWER: B C

Explanation:

CI/CD is characterized by introducing automation into the application and database delivery lifecycle and by automatically validating changes as they are integrated. Continuous integration encourages developers to make small, frequent changes to a shared codebase. Each change can trigger an automated pipeline that builds the relevant artifacts and runs defined tests, giving the team rapid feedback on whether the integrated change remains deployable. Continuous delivery or deployment extends this automation toward release, so that tested artifacts can be consistently promoted through delivery environments using repeatable pipeline stages rather than ad hoc manual processes.

Accordingly, "It introduces automation into all stages of app/database development" captures the core CI/CD practice of automating the build, test, integration, and delivery workflow. "Development changes are automatically tested on pushes, commits, and merges" describes the continuous-integration trigger model: repository events initiate automated checks, helping teams discover integration issues early and maintain a reliable shared branch. Oracle Cloud Infrastructure DevOps provides build pipelines, deployment pipelines, source-code integration, and automated delivery capabilities that support these CI/CD practices. See the [Oracle Cloud Infrastructure DevOps documentation](#) and Oracle's [DevOps overview](#) for the role of automation, continuous integration, testing, and delivery in modern software development.

QUESTION NO: 6

Multicloud strategy eliminates the reliance of a single cloud provider but also provides which two advantages? (Choose all correct answers)

A. Ability to have private clouds and maintain customer data centers

B. Compliance for use of specific geographies for data sovereignty laws

C. Duplication of data and applications

D. Workload utilization of specific services

ANSWER: B D

Explanation:

Compliance for use of specific geographies for data sovereignty laws is an advantage because a multicloud architecture lets an organization select providers and cloud regions that satisfy applicable data-residency, sovereignty, and regulatory requirements. This flexibility is important when data must remain in a particular country or jurisdiction, or when services supporting a regulated workload need to be deployed in an approved location. Oracle explains that regions are geographically localized areas, which enables customers to choose where their cloud resources and data reside.

Workload utilization of specific services is also an advantage because cloud providers have differentiated service portfolios, platform capabilities, regional availability, integrations, and operational characteristics. A multicloud strategy allows teams to place a workload with the provider service best suited to its technical and business requirements while still integrating it with services hosted in another cloud. Oracle's multicloud offerings, including the Oracle Database@Azure collaboration, illustrate this approach by enabling customers to use Oracle database services alongside Azure services within a connected multicloud environment. These choices can help organizations align individual workloads with the most suitable cloud capabilities without making every workload dependent on one provider. See [Oracle Cloud Infrastructure Regions and Availability Domains](#) and [Oracle Database@Azure](#).

QUESTION NO: 7

Which two statements are true when deciding which Oracle Cloud Infrastructure (OCI) region to register an Exadata Cloud@Customer infrastructure in?

- A. Exadata Cloud@Customer is hosted in 4 customer data center so the Exadata infrastructure is not registered in an region.
- B. Consider any business policies or regulations that preclude the use of a particular region. The Exadata Cloud@Customer region cannot be changed after the infrastructure is created.
- C. Consider which availability domain (within the region) to create the Exadata Cloud#Customer infrastructure
- D. Consider the physical proximity of the region you register the infrastructure in to your data center.

ANSWER: B D

Explanation:

Consider any business policies or regulations that preclude the use of a particular region. The Exadata Cloud@Customer region cannot be changed after the infrastructure is created. This is an essential planning consideration because registering the infrastructure associates it with a specific OCI region for management and control-plane services. Organizations must therefore select a region that satisfies data-residency obligations, industry rules, contractual requirements, and internal governance policies before registration. Oracle documentation specifically advises customers to account for business policies and regulations that may prevent use of a given region, and the selected region is not subsequently changed for that infrastructure.

Consider the physical proximity of the region you register the infrastructure in to your data center. Although Exadata Cloud@Customer hardware is installed in the customer data center, it communicates with OCI services in its registered region. Selecting a geographically closer region can reduce network latency and improve the responsiveness of management operations and service interactions. This makes proximity an important operational consideration alongside regulatory requirements. Oracle's Exadata Cloud@Customer guidance identifies both regulatory constraints and physical distance from the data center as factors for choosing the registration region. See [Oracle Exadata Cloud@Customer documentation](#) and [Oracle Cloud Regions](#).

QUESTION NO: 8

What is the primary goal of Oracle Maximum Availability Architecture for customer systems?

- A. Scale out, application availability and protection
- B. Active Replication, Data Protection and continuous availability
- C. Continuous Availability, network security and failover
- D. Data Protection, access and reporting

ANSWER: B

Explanation:

Active Replication, Data Protection and continuous availability best represents the primary goal of Oracle Maximum Availability Architecture (MAA). MAA is Oracle's documented framework of architecture patterns, integrated technologies, and operational best practices for keeping business services available while safeguarding data against failures, corruption, site outages, and planned maintenance. It is not limited to recovery after an outage; it aims to prevent downtime where possible and to minimize service interruption when failures do occur.

Active replication is central to this outcome because replicated database environments can maintain protected copies of data and support rapid role transitions or failover. Oracle technologies commonly used in MAA, such as Oracle Data Guard and Oracle Active Data Guard, provide data synchronization, disaster-recovery capability, and the ability to use standby systems for appropriate read-only workloads. Continuous availability reflects the MAA objective of sustaining application access through resilient system design, redundancy, monitoring, and automated recovery procedures. Together, these capabilities align directly with Oracle's stated focus on optimal availability and data protection for mission-critical systems. See Oracle's [Maximum Availability Architecture overview](#) and the [High Availability Overview](#) for the MAA principles and related technologies.

QUESTION NO: 9

Which three are prerequisites for using the MySQL Database Service? (Choose all correct answers)

- A. You must create a tenancy and be able to sign in to it
- B. You must install the downloaded MySQL version onto Oracle Cloud Infrastructure.
- C. You or your group must be granted the policies described in Mandatory Policies.
- D. You must create a compartment to store your resources.

ANSWER: A C D

Explanation:

Using MySQL Database Service begins with access to an Oracle Cloud Infrastructure tenancy. The tenancy is the fundamental OCI account boundary in which identities, permissions, networking, and service resources are managed, so a user must be able to sign in before provisioning or administering a MySQL DB system. Oracle also requires the user or group to have appropriate IAM policies. These policies authorize actions such as creating, managing, and connecting to MySQL DB systems and related OCI resources; without them, access to the service is not sufficient to perform operational tasks. A compartment is also used to store and organize the MySQL Database Service resources. Compartments provide a logical isolation boundary for resources and allow administrators to apply access policies and governance according to a project, team, or workload. This setup supports controlled provisioning and ongoing resource management within the tenancy. Oracle's MySQL Database Service documentation covers the service setup and DB system lifecycle, while its IAM documentation describes the policy-based authorization model used by OCI. See [Oracle MySQL Database Service Getting Started](#) and [OCI IAM Policies documentation](#).