DUMPSARENA

Oracle Cloud Database 2023 Migration and Integration Professional

Oracle 1z0-1094-23

Version Demo

Total Demo Questions: 11

Total Premium Questions: 117

Buy Premium PDF

https://dumpsarena.co

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

Topic Break Down

Topic	No. of Questions
Topic 1, Oracle Cloud Database Migration	80
Topic 2, Oracle Cloud Database Integration	35
Topic 3, Mix Questions	2
Total	117

QUESTION NO: 1

Which TWO methods are supported by Oracle Zero Downtime Migration (ZDM) for migrating onpremises databases to Autonomous Database? (Choose all correct answers)

A. Logical online or offline migration using a database link

B. Logical online or offline migration using OCI Object Storage

C. Physical online or offline migration using OCI Object Storage

D. Physical online or offline migration using a database link

ANSWER: B C

Explanation:

Oracle Zero Downtime Migration supports both logical and physical migration workflows for eligible on-premises Oracle databases moving to Oracle Autonomous Database, with OCI Object Storage used as the supported staging and transfer medium. A logical migration transfers database objects and data through Oracle logical migration tooling, with ZDM orchestrating the workflow and using Object Storage to hold migration artifacts. Depending on the selected workflow and downtime requirements, ZDM supports both online and offline logical migration approaches.

ZDM also supports physical migration workflows using OCI Object Storage. In a physical workflow, database backup and recovery artifacts are transferred through Object Storage, allowing the target Autonomous Database environment to be created from the source database's physical representation. Online physical migration uses incremental synchronization to reduce the final cutover outage, while offline physical migration completes the transfer while the source is unavailable for the migration window. The applicable workflow depends on source database version, platform, and Autonomous Database service requirements. Oracle documents the supported logical and physical migration methods, including their online and offline variants, in the ZDM migration guide and Autonomous Database migration documentation. See Oracle Zero Downtime Migration documentation and Migrate to Autonomous Database.

QUESTION NO: 2

Which autonomous database types are available on Oracle Cloud? (Choose three.)

A. Autonomous Relation Database System

B. Autonomous Data Warehouse

C. Autonomous Streams

D. Autonomous Transaction Processing

E. Autonomous JSON Database

ANSWER: B D E

Explanation:

Oracle Cloud Infrastructure offers purpose-oriented Autonomous Database workload types that retain the platform's automated operations, including automatic provisioning, tuning, scaling, backup, patching, and security capabilities. **Autonomous Data Warehouse** is the Autonomous Database service optimized for analytics, reporting, data marts, and data-warehouse workloads. **Autonomous Transaction Processing** is optimized for transaction-processing, mixed workloads, and application development; it is also the service type associated with Oracle APEX and can support a broad range of transactional application requirements. **Autonomous JSON Database** is an Autonomous Database offering designed for JSON-centric application development. It provides Oracle Database JSON features and API-oriented JSON document access while benefiting from the autonomous operational model. These are recognized Oracle Cloud Autonomous Database service types, allowing customers to select an operational profile aligned with analytics, transactions, or document-centric development needs. Oracle's Autonomous Database documentation describes the available workload types and the capabilities of each service. See Oracle Autonomous Database overview and Oracle Autonomous JSON Database documentation.

QUESTION NO: 3

Which three Oracle technologies are leveraged by Oracle Zero Downtime Migration Physical Online Workflow? (Choose all correct answers)

A. Oracle Data Guard

B. Recovery Manager

C. SQL*Net

D. Oracle Real Application Testing

E. Streams

ANSWER: A B C

Explanation:

Oracle Zero Downtime Migration (ZDM) Physical Online Migration uses Oracle Data Guard, Recovery Manager, and SQL*Net as core components of its migration architecture. Oracle Data Guard provides the physical standby-based replication mechanism: after the target is instantiated, redo generated on the source is transported and applied to keep the target synchronized until switchover. This enables the source database to remain available for normal workload during almost all of the migration. Recovery Manager (RMAN) is used to create and transfer the initial backup needed to instantiate the target database, including the database files and required recovery information. RMAN's backup, restore, and recovery capabilities establish the target in a state from which Data Guard redo apply can continue synchronization. SQL*Net, now commonly documented as Oracle Net Services, supplies the secure Oracle database connectivity used by ZDM and by the database components to communicate across source, target, and migration-host environments. Together, these technologies support a physical online workflow that limits the outage primarily to the final cutover. Oracle's ZDM documentation describes Physical Online Migration as using RMAN backup and restore with Data Guard for ongoing synchronization. See Oracle Zero Downtime Migration documentation and Oracle Data Guard Concepts and Administration.

QUESTION NO: 4

You have an on-premises non-TDE-encrypted Linux 19c container database (CDB) that is already being regularly backed up to Oracle Cloud Infrastructure (OCI) Object Storage via the Database Backup Cloud Service. You plan on using Recovery Manager (RMAN) to migrate it to Database Cloud Service. What is the valid strategy to migrate?

A. Create a new OCI DB Service instance and use the "instantiate from backup" option in the console. This will automatically take care of the encryption.

B. Create a new OCI DB Service instance, delete the existing database, and use the DBaa$ tooling to configure the DB Cloud Backup Module to point to the existing bucket, and then restore the backup by using the console.

C. Create a new OCI DB Service Instance, copy the backup data from the bucket to the FRA, and then use the console to perform a DB restore.

D. Create a new DB Service instance, delete the existing database, manually install the DB Cloud Backup Module pointing it to the backed where the on-premises backups are located, and perform a complted RMAN RESTORE AS encryption operation.

ANSWER: A

Explanation:

Create a new OCI DB Service instance and use the "instantiate from backup" option in the console. This will automatically take care of the encryption. This is the supported streamlined path when the source database is already protected by RMAN backups written to OCI Object Storage through the Database Backup Cloud Service. During DB Service provisioning, the console's instantiate-from-backup workflow uses the selected backup location and Recovery Manager recovery metadata to create the cloud database from the existing backup set. It is designed for bringing an externally backed-up database into a newly provisioned DB system rather than requiring a separate manual staging or restoration workflow.

The source being a non-TDE-encrypted 19c CDB is important because Oracle Database Cloud Service protects databases at rest using Transparent Data Encryption. The service provisioning and instantiation process manages the encryption requirements for the target cloud database, including the necessary keystore handling, as part of creating the DB Service database from the backup. This allows the migration to use the existing OCI Object Storage backups while producing a database that conforms to the cloud service's encryption model. Oracle documents this workflow under Instantiate a DB System from a Backup; see also Oracle's Backup and Recovery documentation for DB Service.

QUESTION NO: 5

In which database edition does Data Pump support parallelism?

A. Both Standard Edition and Enterprise Edition

B. Enterprise Edition only

C. Standard Edition only

D. Data Pump does not support parallelism.

ANSWER: B

Explanation:

Enterprise Edition only is correct. Oracle Data Pump can execute export and import work through multiple worker processes when the `PARALLEL` parameter is set to a value greater than one. This permits multiple data objects, or partitions and subpartitions where applicable, to be processed concurrently, which can substantially reduce elapsed time for large migration, backup, and loading operations. Oracle documents the Data Pump parallel-execution capability and specifies that parallel Data Pump jobs are available only with Oracle Database Enterprise Edition. The effective degree of parallelism also depends on factors such as the number of dump files, available CPU and I/O capacity, and the nature of the exported objects. For example, an export requires sufficient dump-file capacity for workers to write concurrently, commonly using a dump-file name template containing `%U`. The Data Pump Export documentation describes use of the `PARALLEL` parameter and its relationship to worker processes, while Oracle's licensing documentation identifies Enterprise Edition as the edition supporting parallel Data Pump operations. See the Data Pump Export PARALLEL parameter documentation and the Oracle Database Licensing Information User Manual.

QUESTION NO: 6

You have plugged a non-CDB database into a target container database as a PDB. The PDB is open in upgrade mode and its data dictionary must be converted to the multitenant architecture.

Which script must be run in the new PDB?

A. datapatch

B. catctl.pl

C. utlrp.sql

D. noncdb_to_pdb.sql

ANSWER: D

Explanation:

`noncdb_to_pdb.sql` performs the required conversion of a non-CDB after it has been plugged into a CDB as a PDB. The script updates dictionary metadata so that the database can operate as a normal pluggable database.

`datapatch` applies SQL changes associated with Oracle home patches. `catctl.pl` is used for database upgrade operations, and `utlrp.sql` recompiles invalid objects but does not perform the non-CDB-to-PDB conversion.

QUESTION NO: 7

You are asked to isolate all the erroneous rows into error tables before they get into the target tables, in order to have a list of corrections that need to be addressed before the next run.

At the next run, the demand is to check whether there are some rows that have been fixed and automatically include them in the integration flow.

Which two conditions need to be met to accomplish this demand in Oracle Data Integrator (ODI)? (Choose two.)

A. You must have an update key on the target table.

B. The Check Knowledge Module must have the RECYCLE ERRORS option set to TRUE.

C. The static control has to be enabled and the RECYCLE ERRORS option set to TRUE.

D. The flow control has to be enabled and the RECYCLE ERRORS option set to TRUE.

E. You must have a primary key on the target table.

ANSWER: D E

Explanation:

Corrected error rows can be reintegrated only when ODI is performing flow control and error recycling is enabled. With flow control enabled, ODI validates incoming rows against the target datastore's constraints before the rows are committed to the target. Rows that fail validation are written to ODI's error table, preserving both the rejected data and the information needed to correct it. Setting `RECYCLE ERRORS` to `TRUE` instructs ODI to reconsider rows previously captured as errors on a later execution, so rows that have since been corrected can re-enter the normal integration flow automatically.

A primary key on the target table is also required because it provides the unique row identity ODI needs to manage the target data and reliably relate the controlled/recycled records to the target structure. Together, target-key metadata and enabled flow control with error recycling provide the controlled-load process required to retain rejected records for remediation and subsequently load corrected records. Oracle describes flow control and Knowledge Module behavior in its ODI documentation; see Developing Integration Projects with Oracle Data Integrator and Understanding Oracle Data Integrator.

QUESTION NO: 8

Which TWO methods support a difference in endian format when migrating a database from onpremises to Oracle Cloud? (Choose all correct answers)

A. PDB Unplugging/Plugging

B. Data Pump Conventional Export/Import

C. RMAN cross-Platform Transportable PDB

D. PDB Remote Cloning

ANSWER: B C

Explanation:

Data Pump Conventional Export/Import supports migrations between platforms with different endian formats because it is a logical migration method. The export writes database object definitions and table data into Data Pump dump files rather than copying the database's physical data files. At the target, Import creates the objects and loads the data using the target database format, so no physical data-file endian compatibility is required. This makes it a broadly applicable approach for moving schemas or full databases to an Oracle Cloud database on a platform with a different byte order. Oracle describes Data Pump as the utility for logical export and import of database data and metadata in the Oracle Data Pump documentation.

RMAN cross-Platform Transportable PDB also supports a different endian format. Oracle RMAN can transport a pluggable database across platforms and perform the required conversion of the PDB data files when the source and destination have different endian formats. This preserves the transportable-PDB workflow while making the physical files usable by the destination platform. The procedure requires the applicable cross-platform and database-version prerequisites, but it is specifically intended to handle cross-endian PDB movement. Oracle documents RMAN transport and conversion support, including cross-platform database-file transport, in its RMAN Transporting Data documentation.

QUESTION NO: 9

Which type of Zero Downtime Migration (ZDM) can you use for Inflight Upgrade?

A. Both Physical and Logical Migration is supported.

B. ZDM does not support Inflight Upgrade.

C. Logical Migration only

D. Physical Migration only

ANSWER: C

Explanation:

Logical Migration only is correct because ZDM's in-flight upgrade capability is implemented through the logical migration workflow. This workflow uses Oracle GoldenGate replication to keep the target synchronized while the source database remains available, enabling the target to be created at a later Oracle Database release than the source. ZDM performs the initial logical load, establishes ongoing change capture and replication, and then coordinates a controlled switchover once the target has caught up. This makes it possible to combine a migration with a database-version upgrade while minimizing application downtime. In practical terms, the logical approach is appropriate because schema metadata and data are moved into a target database that can be configured for the desired target release, rather than requiring a physical copy of the source database's existing database version and file structure. Oracle's Zero Downtime Migration documentation describes logical online migration as the Oracle GoldenGate-based migration path and documents the supported source-to-target version combinations and upgrade scenarios. See the Oracle Zero Downtime Migration documentation and the logical migration workflow guide.

QUESTION NO: 10

Which platform and version supports the deployment of Oracle Zero Downtime Migration Service host?

A. Oracle Linux 7

B. Windows 2012 R2 and 2016 R1

C. IBMAIX7.1 and later WW

D. Red Hat Enterprise Linux 8.2

E. Solaris 11 Update 3 and later

ANSWER: A

Explanation:

Oracle Linux 7 is the supported deployment platform identified for the Oracle Zero Downtime Migration (ZDM) service host in this question's ZDM platform requirements. The ZDM service host is the centrally managed system on which ZDM software runs and from which migration workflows are orchestrated. It must use a supported Linux operating-system environment so that the required Oracle software, utilities, SSH connectivity, database client components, and migration tooling can be installed and managed reliably.

Oracle Linux is Oracle's enterprise Linux distribution and is designed and tested for Oracle software workloads. Deploying the service host on Oracle Linux 7 provides the supported base environment for configuring ZDM prerequisites, registering

source and target databases, and executing physical or logical migration jobs while ZDM coordinates the migration process. Oracle's ZDM documentation describes the service host preparation process and specifies the operating-system requirements applicable to the ZDM release. See Preparing the Zero Downtime Migration Service Host and Oracle's Oracle Linux overview for the platform context.

QUESTION NO: 11

You are trying to load a denormalized target table C and the source attributes come from table A and Table B in Oracle Data Integrator (ODI). Which statement is true?

A. A mapping can have only one source and one target. You need to create 2 mappings to load target C.

B. You can join table A and table B, as required by the data model and load target C in a Single mapping

C. You are required to do inner join between table A and table B in order to load target C.

D. You are required to do outer join between table A and table B in order to load target C.

ANSWER: B

Explanation:

You can join table A and table B, as required by the data model and load target C in a Single mapping. An Oracle Data Integrator mapping is designed to define the flow of data from one or more datastores to one or more targets. When a denormalized target needs attributes held in separate source tables, both source datastores can be added to the same logical mapping. A join can then be defined between their attributes, and the resulting combined rows can be mapped to the columns of target table C. ODI generates and executes the appropriate integration code using the selected logical and physical design, including the knowledge modules configured for the source and target technologies. The appropriate join type is determined by the intended business result and the source data relationships; ODI does not impose a single mandatory join type for this scenario. Keeping the join and target load in one mapping provides a clear, maintainable representation of the transformation and enables ODI to optimize the generated flow according to the configured staging area and knowledge modules. Oracle describes mappings as the design-time objects used to map source data to target datastores and perform transformations, including joins. See the Oracle Data Integrator Developer's Guide: Developing Integration Projects and the Oracle documentation on creating and using mappings.