

1Z0-931-25 Deutsch Prüfung, 1Z0-931-25 Exam Fragen



Außerdem sind jetzt einige Teile dieser EchteFrage 1Z0-931-25 Prüfungsfragen kostenlos erhältlich: <https://drive.google.com/open?id=1nW3Q14X389LGSHLIYJaAO4XMLuRyleBK>

Unser EchteFrage ist eine fachliche IT-Website. Ihre Erfolgsquote beträgt 100%. Viele Kandidaten haben das schon bewiesen. Weil wir ein riesiges IT-Expertenteam hat, das nach ihren fachlichen Erfahrungen und Kenntnissen die Oracle 1Z0-931-25 Prüfungsfragen und Antworten bearbeitet, um die Interessen der Kandidaten zu schützen und zugleich ihren Bedürfnisse abzudecken. Nach den Bedürfnissen der Kandidaten haben sie zielgerichtete und anwendbare Schulungsmaterialien entworfen, nämlich die Schulungsunterlagen zur Oracle 1Z0-931-25 Zertifizierungsprüfung, die Fragen und Antworten enthalten.

Oracle 1Z0-931-25 Prüfungsplan:

Thema	Einzelheiten
Thema 1	Autonomous Database Shared: This section of the exam measures the skills of Cloud Engineers and focuses on creating and managing shared Autonomous Database instances. It includes provisioning, scaling, and starting or stopping instances, as well as database consolidation with Elastic Resource Pools. It also covers user management, cloning, database migration, monitoring, backup and restore processes, and introduces Data Guard for high availability, ensuring cloud engineers can maintain optimal database performance.
Thema 2	Developing on Autonomous Database: This section of the exam measures the skills of Application Developers and focuses on developing and extending applications using Autonomous Database. It covers using generative AI for natural language queries, Autonomous JSON Database, Oracle Text for document search, location-based analysis with Autonomous Spatial, Autonomous Graph for data relationships, and integration with Object Storage, enabling developers to build intelligent, scalable applications.
Thema 3	Data Lake Analytics with Autonomous Database: This section of the exam measures the skills of Big Data Engineers and explores how Autonomous Database can be used for analytics in data lake environments. It includes data ingestion, query optimization, and leveraging cloud-native analytics services, ensuring engineers can efficiently process and analyze large volumes of structured and unstructured data.

Thema 4	Autonomous Database Tools: This section of the exam measures the skills of Data Analysts and covers the tools available within Autonomous Databases for advanced data processing and analytics. It includes Oracle Machine Learning, APEX, and SQL Developer Web for database development, as well as data transformation, business model creation, data insights, and data analysis, allowing analysts to extract valuable insights from large datasets.
Thema 5	Migrating to Autonomous Database: This section of the exam measures the skills of Cloud Migration Specialists and covers strategies for migrating existing databases to Autonomous Database. It includes understanding migration considerations, and available options, and using Oracle Data Pump to transfer data seamlessly while minimizing downtime, ensuring smooth transitions to Oracle Cloud infrastructure.
Thema 6	Managing and Maintaining Autonomous Database: This section of the exam measures the skills of Database Administrators and focuses on the ongoing management and maintenance of Autonomous Database instances. It includes using REST APIs and OCI CLI for automation, configuring access control lists and private endpoints, monitoring performance, setting up notifications, utilizing features like auto-indexing and data safe, handling connectivity through wallets and service handles, and configuring disaster recovery using Data Guard to ensure business continuity.

>> 1Z0-931-25 Deutsch Prüfung <<

Kostenlos 1Z0-931-25 Dumps Torrent & 1Z0-931-25 exams4sure pdf & Oracle 1Z0-931-25 pdf vce

EchteFrage ist eine Schulungswebsite, die spezielle Fragen und Antworten zur Oracle 1Z0-931-25 Zertifizierungsprüfung IT-Zertifizierungsprüfung und Prüfungsthemen bieten. Gegen die populäre Oracle 1Z0-931-25 Zertifizierungsprüfung haben wir neuen Schulungskonzepte entwickelt, die die Bedürfnisse vieler Leute abdecken können. Viele berühmten IT-Firmen stellen ihre Angestellte laut dem Oracle 1Z0-931-25 Zertifikat ein. Deahalb ist die Oracle 1Z0-931-25 (Oracle Autonomous Database Cloud 2025 Professional) Zertifizierungsprüfung zur Zeit sehr populär. EchteFrage wird von vielen akzeptiert und hat den Traum einer Mehrheit der Leute verwirklicht. Wenn Sie mit Hilfe von EchteFrage die Prüfung nicht bestehen, zahlen wir Ihnen die gesamte Summe zurück.

Oracle Autonomous Database Cloud 2025 Professional 1Z0-931-25 Prüfungsfragen mit Lösungen (Q112-Q117):

112. Frage

In which four ways can Oracle Database optimally access data in Object Storage? (Choose four.)

- A. Optimized data archive using partitioned external tables
- B. Optimized data archive using hybrid partitioned tables
- C. Scan avoidance using block skipping when reading parquet and orc files
- D. Scan avoidance using columnar pruning for .csv files
- E. Scan avoidance using partitioned external tables
- F. Scan avoidance using columnar pruning for columnar stores like parquet and orc

Antwort: A,B,E,F

Begründung:

Oracle Database provides several techniques to optimize data access from Object Storage, particularly in the context of Autonomous Database, leveraging external tables and advanced storage formats. The question asks for four correct methods, and based on Oracle documentation, the following are the most applicable:

Correct Answer (A): Scan avoidance using partitioned external tables

Partitioned external tables allow Oracle Database to skip irrelevant partitions when querying data stored in Object Storage. By organizing data into partitions (e.g., by date or region), the database engine can prune partitions that don't match the query predicates, significantly reducing the amount of data scanned and improving performance. This is a well-documented optimization for external data access in Oracle Database and Autonomous Database environments.

Correct Answer (D): Scan avoidance using columnar pruning for columnar stores like parquet and orc Columnar pruning is a technique where only the required columns are read from columnar file formats such as Parquet or ORC stored in Object Storage.

These formats store data column-wise, enabling the database to avoid scanning entire rows or irrelevant columns, which is particularly efficient for analytical queries common in Autonomous Data Warehouse (ADW). This is a standard optimization supported by Oracle's external table framework when accessing Object Storage.

Correct Answer (E): Optimized data archive using hybrid partitioned tables Hybrid partitioned tables combine local database partitions with external partitions stored in Object Storage. This allows older, less frequently accessed data to be archived efficiently in the cloud while remaining queryable alongside active data in the database. The database optimizes access by seamlessly integrating these partitions, reducing costs and improving archival efficiency. This feature is explicitly supported in Oracle Database and enhanced in Autonomous Database for data lifecycle management.

Correct Answer (F): Optimized data archive using partitioned external tables Similar to hybrid partitioned tables, using partitioned external tables alone optimizes data archiving by storing historical data in Object Storage with partitioning (e.g., by year). This method enables efficient querying of archived data by pruning unneeded partitions, offering a cost-effective and scalable archival solution. It's a distinct approach from hybrid tables, focusing solely on external storage, and is widely used in Oracle environments.

Incorrect Options:

B . Scan avoidance using columnar pruning for .csv files

CSV files are row-based, not columnar, and lack the internal structure of formats like Parquet or ORC. While Oracle can read CSVs from Object Storage via external tables, columnar pruning is not applicable because CSVs don't support column-wise storage or metadata for pruning. This makes this option incorrect as a specific optimization technique, though basic predicate pushdown might still reduce scanning to some extent.

C . Scan avoidance using block skipping when reading parquet and orc files Block skipping (or row group skipping) is a feature in some database systems where metadata in Parquet or ORC files allows skipping entire blocks of data based on query filters. While Oracle supports Parquet and ORC through external tables and can leverage their columnar nature (via pruning), "block skipping" is not explicitly highlighted as a primary optimization in Oracle's documentation for Autonomous Database. It's more commonly associated with systems like Apache Spark or Hive. Oracle's focus is on columnar pruning and partitioning, making this option less accurate in this context.

Why Four Answers?

The question specifies "four ways," and while six options are provided, A, D, E, and F are the most directly supported and documented methods in Oracle Autonomous Database for optimizing Object Storage access. Options B and C, while conceptually related to data access optimizations, are either inapplicable (CSV lacks columnar structure) or not explicitly emphasized (block skipping) in Oracle's feature set for this purpose.

This selection aligns with Oracle's focus on partitioning and columnar formats for efficient cloud data access, ensuring both performance and archival optimization.

Reference:

External Tables and Object Storage

Hybrid Partitioned Tables

Autonomous Database Data Loading

113. Frage

Which statement is FALSE about Data Insights?

- A. Data Insights provides a wide range of graphical data presentation capabilities
- **B. Data Insights are automatically generated by various analytic functions built into the database**
- C. The results of the Insight analysis appear as a series of bar charts in the Data Insights dashboard
- D. Data Insights display information about patterns and anomalies in the data of entities in your Oracle Autonomous Database

Antwort: B

Begründung:

Data Insights is a feature in Autonomous Database that helps users understand their data. The false statement is:

Data Insights are automatically generated by various analytic functions built into the database (C): This is incorrect. Data Insights are not solely the result of automatic execution of built-in analytic functions (e.g., AVG, SUM, or RANK). Instead, they are generated through a combination of user-initiated analysis and Oracle's machine learning-driven capabilities within the Data Insights dashboard (part of Database Actions or OCI console). Users select datasets or tables, and the system applies algorithms to identify patterns (e.g., trends in sales) or anomalies (e.g., outlier transactions), but this process isn't just a passive outcome of pre-existing database functions—it's an active, curated feature requiring configuration. For example, a user might explore a SALES table, and Data Insights highlights a spike in Q4 sales, but this requires user input to define scope, not just automatic function output.

The true statements are:

Data Insights display information about patterns and anomalies in the data of entities in your Oracle Autonomous Database (A):

True. The feature visualizes trends (e.g., seasonal sales increases) and outliers (e.g., unexpected data drops) in tables or views, helping users spot significant data behaviors. For instance, it might show a bar chart of monthly revenue with an anomaly flagged for a sudden dip.

Data Insights provides a wide range of graphical data presentation capabilities (B): True. It offers visualizations like bar charts, line graphs, and scatter plots, customizable to represent data insights effectively. E.g., a line graph might track customer sign-ups over time, with options to adjust axes or filters.

The results of the Insight analysis appear as a series of bar charts in the Data Insights dashboard (D): True, partially. While bar charts are a common default (e.g., comparing sales by region), the dashboard supports multiple chart types, but the statement's focus on bar charts aligns with typical output for simple insights.

The misconception in C overlooks the interactive, ML-assisted nature of Data Insights, distinguishing it from passive function-based analytics.

114. Frage

A corporation is building a web application to allow its customers to schedule service requests online. There is also a need to run operational reports at times during non-peak hours. The architecture team is debating whether such reports should be run on the OLTP database or in a separate data mart. The DBA Manager does not want to add any more admin responsibility to the team and is looking for a database option that's low to zero maintenance, but meets their strict performance requirements as well. Which Oracle Cloud Infrastructure database service is appropriate for this scenario?

- A. ATP using 'tpurgent' and 'high' TNS services to separate connection types
- B. Since the application needs to be highly available, it should be deployed on a Kubernetes Cluster
- C. ADW since operational reporting is a higher priority in this scenario
- D. It is best to build a separate data warehouse, and move the OLTP data on a nightly basis

Antwort: A

Begründung:

The scenario requires a low-maintenance, high-performance database for an OLTP web application with occasional reporting. The correct answer is:

ATP using 'tpurgent' and 'high' TNS services to separate connection types (A): Autonomous Transaction Processing (ATP) is ideal here. It's a fully managed database optimized for OLTP workloads (e.g., scheduling service requests) with zero maintenance overhead-Oracle handles patching, backups, and tuning. ATP supports multiple connection services:

'tpurgent': Prioritizes low-latency, time-critical transactions (e.g., customer scheduling requests), ensuring fast response times for the web app.

'high': Suited for high-concurrency or reporting queries, allowing operational reports to run during non-peak hours without impacting the OLTP workload.

By using these predefined services, the architecture separates transactional and reporting workloads within the same database, meeting performance needs without additional administration. For example, the web app connects via tpurgent for real-time updates, while a reporting tool uses high for batch queries at night, leveraging ATP's auto-scaling if needed.

The incorrect options are:

Since the application needs to be highly available, it should be deployed on a Kubernetes Cluster (B): This misinterprets the question-it's about the database service, not the application deployment. Kubernetes is for container orchestration, not a database solution, and adds complexity counter to the low-maintenance goal.

It is best to build a separate data warehouse, and move the OLTP data on a nightly basis (C): While a separate Autonomous Data Warehouse (ADW) could handle reporting, it requires data movement (e.g., via ETL), increasing admin effort and complexity, which the DBA Manager wants to avoid. ATP can handle both workloads with proper service separation.

ADW since operational reporting is a higher priority in this scenario (D): ADW is optimized for analytics, not OLTP. The web app's transactional needs are primary, with reporting secondary and occasional, making ATP more suitable.

ATP's self-managing nature and service flexibility make it the best fit.

115. Frage

Which three statements are true about procedures in the DBMS_CLOUD package?

- A. The DBMS_CLOUD.VALIDATE_EXTERNAL_TABLE validates the source file for an external table, generates log information, and stores the rows that do not match the format options specified for the external table in a bad file table in Autonomous Data Warehouse.
- B. The DBMS_CLOUD.CREATE_CREDENTIAL procedure stores cloud object credentials in the Autonomous Data Warehouse.
- C. The DBMS_CLOUD.PUT_OBJECT procedure copies a file from cloud object storage to the Autonomous Data Warehouse.
- D. The DBMS_CLOUD.CREATE_EXTERNAL_TABLE procedure creates an external table on files in the cloud. You can

run queries on external data from the Autonomous Data Warehouse.

- E. The DBMS_CLOUD.DELETE_FILE procedure removes the credential file from the Autonomous Data Warehouse.

Antwort: A,B,D

Begründung:

Full Detailed In-Depth Explanation:

The DBMS_CLOUD package facilitates cloud integration:

A: True. CREATE_EXTERNAL_TABLE creates external tables for querying cloud data directly.

B: False. No DELETE_FILE procedure exists; credential management uses DROP_CREDENTIAL.

C: True. CREATE_CREDENTIAL securely stores cloud storage credentials.

D: True. VALIDATE_EXTERNAL_TABLE checks file integrity and logs errors.

E: False. PUT_OBJECT uploads files to cloud storage, not into the database.

116. Frage

Which three operations are available for a quarterly maintenance update on Autonomous Container Database? (Choose three.)

- A. Skip patching
- B. Reschedule and patch immediately
- C. Apply a partial patch
- D. Patch immediately
- E. Reschedule patching

Antwort: A,B,E

Begründung:

Oracle Autonomous Container Database (ACD) provides specific options for managing quarterly maintenance updates to ensure minimal disruption and flexibility. The three correct operations available are:

Reschedule and patch immediately (B): This option allows users to adjust the maintenance window to a more convenient time and apply the patch right away within that rescheduled window. It's useful when the default schedule conflicts with business operations but immediate patching is still desired.

Reschedule patching (C): This option enables users to defer the patching to a later maintenance window that suits their operational needs, providing flexibility without applying the patch immediately.

Skip patching (D): Users can choose to skip a quarterly maintenance update entirely, which is beneficial if the current database version is stable and no immediate updates are required. Oracle allows skipping up to two consecutive quarterly updates for Autonomous Container Databases on dedicated infrastructure.

The incorrect options are:

Patch immediately (A): This is not a standalone option for quarterly updates on ACDs. Patching occurs within scheduled maintenance windows, and immediate patching outside of rescheduling is not supported as a distinct choice.

Apply a partial patch (E): Oracle does not support applying partial patches during quarterly maintenance updates. Updates are delivered as complete bundles to ensure consistency and security.

This aligns with Oracle's maintenance policies for Autonomous Databases on dedicated infrastructure, where control over scheduling and skipping is provided, but partial patching is not an option.

117. Frage

.....

Die Oracle 1Z0-931-25 Zertifizierungsprüfung ist eine Prüfung, die IT-Technik testet. EchteFrage ist eine Website, die Ihnen zum Bestehen der Oracle 1Z0-931-25 Zertifizierungsprüfung verhilft. Viele Menschen verwenden viel Zeit und Energie auf die Oracle 1Z0-931-25 Zertifizierungsprüfung oder sie geben viel Geld für die Kurse aus, um die Oracle 1Z0-931-25 Zertifizierungsprüfung zu bestehen. Mit EchteFrage brauchen Sie nicht so viel Geld, Zeit und Energie. Die zielgerichteten Übungen von EchteFrage dauern nur 20 Stunden. Sie können dann die Oracle 1Z0-931-25 Zertifizierungsprüfung leicht bestehen.

1Z0-931-25 Exam Fragen: <https://www.echtefrage.top/1Z0-931-25-deutsch-pruefungen.html>

- 1Z0-931-25 Probesfragen □ 1Z0-931-25 Exam Fragen □ 1Z0-931-25 Prüfungsübungen □ Öffnen Sie die Website ▷ www.deutschpruefung.com ◁ Suchen Sie 「 1Z0-931-25 」 Kostenloser Download □ 1Z0-931-25 Probesfragen
- 1Z0-931-25 Schulungsmaterialien - 1Z0-931-25 Dumps Prüfung - 1Z0-931-25 Studienguide □ Suchen Sie jetzt auf ✨ www.itzert.com □ ✨ □ nach □ 1Z0-931-25 □ und laden Sie es kostenlos herunter □ 1Z0-931-25 Originale Fragen

- 2025 Die neuesten EchteFrage 1Z0-931-25 PDF-Versionen Prüfungsfragen und 1Z0-931-25 Fragen und Antworten sind kostenlos verfügbar: <https://drive.google.com/open?id=1nW3Q14X389LGSHLIYJaAO4XMLuRyleBK>

2025 Die neuesten EchteFrage 1Z0-931-25 PDF-Versionen Prüfungsfragen und 1Z0-931-25 Fragen und Antworten sind kostenlos verfügbar: <https://drive.google.com/open?id=1nW3Q14X389LGSHLIYJaAO4XMLuRyleBK>