

ARA-C01 echter Test & ARA-C01 sicherlich-zu-bestehen & ARA-C01 Testguide



2026 Die neuesten ITZert ARA-C01 PDF-Versionen Prüfungsfragen und ARA-C01 Fragen und Antworten sind kostenlos verfügbar: <https://drive.google.com/open?id=1bAvjGte2fomVdxilzZFkuldwKMb2Zscj>

Es ist unnötig für Sie, viel Zeit an einer ARA-C01 Zertifizierungsprüfung zu verwenden. Wenn Sie es schwierig für die Vorbereitung der Snowflake ARA-C01 Prüfung finden und viel Zeit verschwenden müssen, sollen Sie am Besten ITZert ARA-C01 Dumps als Ihr Lerngerät benutzen, weil es kann viel Zeit für Sie sparen. Und es ist wichtiger, dass sie Ihnen versprechen, die Snowflake ARA-C01 Prüfung zu bestehen. Und es gibt keine anderen Unterlagen in dem Markt. Sie können viele andere interessante Dinge machen, statt die Snowflake ARA-C01 Prüfungen vorzubereiten. So, klicken Sie ITZert Webseite und Informieren Sie sich. Sie werden bereuen, diese Chance zu verlieren.

Die Snowflake ARA-C01 (SnowPro Advanced Architect Certification) Prüfung ist eine anspruchsvolle und umfassende Prüfung, die fortgeschrittene Kenntnisse und Fähigkeiten in der Snowflake-Datenplattform testet. Die Prüfung richtet sich an Personen, die ein tiefes Verständnis der Snowflake-Architektur und ihrer fortgeschrittenen Funktionen haben, einschließlich Datenmodellierung, Leistungsabstimmung, Sicherheit und Governance.

Die Snowflake ARA-C01-Prüfung testet das Wissen des Kandidaten in mehreren Bereichen, einschließlich der Snowflake-Architektur, der Datenmodellierung, der Leistungsoptimierung, der Sicherheit und der Governance. Es ist für Fachleute konzipiert, die über umfangreiche Erfahrung in der Arbeit mit Snowflake verfügen und ein tiefes Verständnis für seine Merkmale und Fähigkeiten haben. Die Prüfung ist eine Herausforderung, und das Bestehen erfordert ein gründliches Verständnis der fortschrittlichen Funktionen und Best Practices von Snowflake.

Neueste ARA-C01 Pass Guide & neue Prüfung ARA-C01 braindumps & 100% Erfolgsquote

Es ist nicht so einfach, die ARA-C01 Prüfung zu bestehen. ARA-C01 Prüfung erfordert ein hohes Maß an Fachwissen der IT. Wenn es Ihnen dieses Wissen fehlt, kann ITZert Ihnen die Kenntnissequellen zur Verfügung stehen. Mit ihren reichen Fachkenntnissen und Erfahrungen bietet der Expertenteam die relevanten Fragen und Antworten der ARA-C01 Zertifizierungsprüfung. Wenn Sie ITZert wählen, versprechen wir Ihnen nicht nur eine 100%-Pass-Garantie, sondern stellt Ihnen auch einen einjährigen kostenlosen Update-Service zur Verfügung. Falls Sie in der Prüfung durchfallen, zahlen wir Ihnen die gesamte Summe zurück.

Snowflake SnowPro Advanced Architect Certification ARA-C01 Prüfungsfragen mit Lösungen (Q79-Q84):

79. Frage

Which command will create a schema without Fail-safe and will restrict object owners from passing on access to other users?

- A. create schema EDW.ACCOUNTING WITH MANAGED ACCESS DATA_RETENTION_TIME_IN_DAYS - 7;
- **B. create TRANSIENT schema EDW.ACCOUNTING WITH MANAGED ACCESS DATA_RETENTION_TIME_IN_DAYS = 7;**
- C. create schema EDW.ACCOUNTING WITH MANAGED ACCESS;
- D. create TRANSIENT schema EDW.ACCOUNTING WITH MANAGED ACCESS DATA_RETENTION_TIME_IN_DAYS = 1;

Antwort: B

Begründung:

A transient schema in Snowflake is designed without a Fail-safe period, meaning it does not incur additional storage costs once it leaves Time Travel, and it is not protected by Fail-safe in the event of a data loss. The WITH MANAGED ACCESS option ensures that all privilege grants, including future grants on objects within the schema, are managed by the schema owner, thus restricting object owners from passing on access to other users¹.

Reference =

- * Snowflake Documentation on creating schemas¹
- * Snowflake Documentation on configuring access control²
- * Snowflake Documentation on understanding and viewing Fail-safe³

80. Frage

A large manufacturing company runs a dozen individual Snowflake accounts across its business divisions.

The company wants to increase the level of data sharing to support supply chain optimizations and increase its purchasing leverage with multiple vendors.

The company's Snowflake Architects need to design a solution that would allow the business divisions to decide what to share, while minimizing the level of effort spent on configuration and management. Most of the company divisions use Snowflake accounts in the same cloud deployments with a few exceptions for European-based divisions.

According to Snowflake recommended best practice, how should these requirements be met?

- A. Migrate the European accounts in the global region and manage shares in a connected graph architecture. Deploy a Data Exchange.
- B. Deploy a Private Data Exchange in combination with data shares for the European accounts.
- C. Deploy to the Snowflake Marketplace making sure that `invoker_share()` is used in all secure views.
- **D. Deploy a Private Data Exchange and use replication to allow European data shares in the Exchange.**

Antwort: D

Begründung:

According to Snowflake recommended best practice, the requirements of the large manufacturing company should be met by deploying a Private Data Exchange in combination with data shares for the European accounts. A Private Data Exchange is a feature of the Snowflake Data Cloud platform that enables secure and governed sharing of data between organizations. It allows Snowflake

customers to create their own data hub and invite other parts of their organization or external partners to access and contribute data sets. A Private Data Exchange provides centralized management, granular access control, and data usage metrics for the data shared in the exchange¹. A data share is a secure and direct way of sharing data between Snowflake accounts without having to copy or move the data. A data share allows the data provider to grant privileges on selected objects in their account to one or more data consumers in other accounts². By using a Private Data Exchange in combination with data shares, the company can achieve the following benefits:

- * The business divisions can decide what data to share and publish it to the Private Data Exchange, where it can be discovered and accessed by other members of the exchange. This reduces the effort and complexity of managing multiple data sharing relationships and configurations.
- * The company can leverage the existing Snowflake accounts in the same cloud deployments to create the Private Data Exchange and invite the members to join. This minimizes the migration and setup costs and leverages the existing Snowflake features and security.
- * The company can use data shares to share data with the European accounts that are in different regions or cloud platforms. This allows the company to comply with the regional and regulatory requirements for data sovereignty and privacy, while still enabling data collaboration across the organization.
- * The company can use the Snowflake Data Cloud platform to perform data analysis and transformation on the shared data, as well as integrate with other data sources and applications. This enables the company to optimize its supply chain and increase its purchasing leverage with multiple vendors.

81. Frage

How can the Snowpipe REST API be used to keep a log of data load history?

- **A. Call loadHistoryScan every 10 minutes for a 15-minute time range.**
- B. Call insertReport every 8 minutes for a 10-minute time range.
- C. Call insertReport every 20 minutes, fetching the last 10,000 entries.
- D. Call loadHistoryScan every minute for the maximum time range.

Antwort: A

Begründung:

Snowpipe is a service that automates and optimizes the loading of data from external stages into Snowflake tables. Snowpipe uses a queue to ingest files as they become available in the stage. Snowpipe also provides REST endpoints to load data and retrieve load history reports¹.

The loadHistoryScan endpoint returns the history of files that have been ingested by Snowpipe within a specified time range. The endpoint accepts the following parameters²:

pipe: The fully-qualified name of the pipe to query.

startTimeInclusive: The start of the time range to query, in ISO 8601 format. The value must be within the past 14 days.

endTimeExclusive: The end of the time range to query, in ISO 8601 format. The value must be later than the start time and within the past 14 days.

recentFirst: A boolean flag that indicates whether to return the most recent files first or last. The default value is false, which means the oldest files are returned first.

showSkippedFiles: A boolean flag that indicates whether to include files that were skipped by Snowpipe in the response. The default value is false, which means only files that were loaded are returned.

The loadHistoryScan endpoint can be used to keep a log of data load history by calling it periodically with a suitable time range. The best option among the choices is D, which is to call loadHistoryScan every 10 minutes for a 15-minute time range. This option ensures that the endpoint is called frequently enough to capture the latest files that have been ingested, and that the time range is wide enough to avoid missing any files that may have been delayed or retried by Snowpipe. The other options are either too infrequent, too narrow, or use the wrong endpoint³.

Reference:

1: Introduction to Snowpipe | Snowflake Documentation

2: loadHistoryScan | Snowflake Documentation

3: Monitoring Snowpipe Load History | Snowflake Documentation

82. Frage

A global company needs to securely share its sales and inventory data with a vendor using a Snowflake account.

The company has its Snowflake account in the AWS eu-west-2 Europe (London) region. The vendor's Snowflake account is on the Azure platform in the West Europe region. How should the company's Architect configure the data share?

- A. 1. Promote an existing database in the company's local account to primary.
2. Replicate the database to Snowflake on Azure in the West-Europe region.
3. Create a share and add objects to the share.
4. Add a consumer account to the share for the vendor to access.
- B. 1. Create a new role called db_share.
2. Grant the db_share role privileges to read data from the company database and schema.
3. Create a user for the vendor.
4. Grant the ds_share role to the vendor's users.
- C. 1. Create a share.
2. Add objects to the share.
3. Add a consumer account to the share for the vendor to access.
- D. 1. Create a share.
2. Create a reader account for the vendor to use.
3. Add the reader account to the share.

Antwort: C

Begründung:

The correct way to securely share data with a vendor using a Snowflake account on a different cloud platform and region is to create a share, add objects to the share, and add a consumer account to the share for the vendor to access. This way, the company can control what data is shared, who can access it, and how long the share is valid. The vendor can then query the shared data without copying or moving it to their own account. The other options are either incorrect or inefficient, as they involve creating unnecessary reader accounts, users, roles, or database replication.

<https://learn.snowflake.com/en/certifications/snowpro-advanced-architect/>

83. Frage

An Architect needs to grant a group of ORDER_ADMIN users the ability to clean old data in an ORDERS table (deleting all records older than 5 years), without granting any privileges on the table. The group's manager (ORDER_MANAGER) has full DELETE privileges on the table.

How can the ORDER_ADMIN role be enabled to perform this data cleanup, without needing the DELETE privilege held by the ORDER_MANAGER role?

- A. This scenario would actually not be possible in Snowflake - any user performing a DELETE on a table requires the DELETE privilege to be granted to the role they are using.
- B. Create a stored procedure that can be run using both caller's and owner's rights (allowing the user to specify which rights are used during execution), and grant USAGE on this procedure to ORDER_ADMIN. The ORDER_MANAGER role owns the procedure.
- C. Create a stored procedure that runs with owner's rights, including the appropriate "> 5 years" business logic, and grant USAGE on this procedure to ORDER_ADMIN. The ORDER_MANAGER role owns the procedure.
- D. Create a stored procedure that runs with caller's rights, including the appropriate "> 5 years" business logic, and grant USAGE on this procedure to ORDER_ADMIN. The ORDER_MANAGER role owns the procedure.

Antwort: C

Begründung:

This is the correct answer because it allows the ORDER_ADMIN role to perform the data cleanup without needing the DELETE privilege on the ORDERS table. A stored procedure is a feature that allows scheduling and executing SQL statements or stored procedures in Snowflake. A stored procedure can run with either the caller's rights or the owner's rights. A caller's rights stored procedure runs with the privileges of the role that called the stored procedure, while an owner's rights stored procedure runs with the privileges of the role that created the stored procedure. By creating a stored procedure that runs with owner's rights, the ORDER_MANAGER role can delegate the specific task of deleting old data to the ORDER_ADMIN role, without granting the ORDER_ADMIN role more general privileges on the ORDERS table. The stored procedure must include the appropriate business logic to delete only the records older than 5 years, and the ORDER_MANAGER role must grant the USAGE privilege on the stored procedure to the ORDER_ADMIN role. The ORDER_ADMIN role can then execute the stored procedure to perform the data cleanup.

References:

* Snowflake Documentation: Stored Procedures

* Snowflake Documentation: Understanding Caller's Rights and Owner's Rights Stored Procedures

• • • • •

ARA-C01 Prüfungs: https://www.itcert.com/ARA-C01_valid-braindumps.html

- P.S. Kostenlose 2026 Snowflake ARA-C01 Prüfungsfragen sind auf Google Drive freigegeben von ITZert verfügbar: <https://drive.google.com/open?id=1bAvjGte2fomVdxilzZFkuldwKMb2Zscj>