

# Professional-Data-Engineer PrüfungGuide, Google Professional-Data-Engineer Zertifikat - Google Certified Professional Data Engineer Exam



Übrigens, Sie können die vollständige Version der ZertFragen Professional-Data-Engineer Prüfungsfragen aus dem Cloud-Speicher herunterladen: [https://drive.google.com/open?id=1wmHqxRJhud83KTHZ1i\\_ELkfPYL7zUHD5](https://drive.google.com/open?id=1wmHqxRJhud83KTHZ1i_ELkfPYL7zUHD5)

ZertFragen ist eine spezielle Website, die Schulungsunterlagen zur Google Professional-Data-Engineer Zertifizierungsprüfung bietet. Hier werden Ihre Fachkenntnisse nicht nur befördert werden. Und Sie können auch die Prüfung einmalig bestehen. Die Schulungsunterlagen von ZertFragen werden von den erfahrungsreichen Fachleuten nach ihren Erfahrungen und Kenntnissen bearbeitet. Sie sind von guter Qualität und extrem genau. ZertFragen wird Ihnen helfen, nicht nur die Google Professional-Data-Engineer Zertifizierungsprüfung zu bestehen, sondern auch Ihre Fachkenntnisse zu konsolidieren. Außerdem genießen Sie einen einjährigen Update-Service.

Die Prüfungsfragen und Antworten zur Google Professional-Data-Engineer Zertifizierung von ZertFragen enthalten unbeschränkte Antwortenspeicherungen. So können Sie ganz mühelos die Prüfung bestehen. Die Schulungsunterlagen zur Google Professional-Data-Engineer Prüfung von ZertFragen sind die besten. Mit deren Hilfe können Sie ganz einfach die Prüfung bestehen und das Zertifikat für Google Professional-Data-Engineer Prüfung erhalten.

>> Professional-Data-Engineer Zertifizierungsantworten <<

## Valid Professional-Data-Engineer exam materials offer you accurate preparation dumps

Es ist uns allen bekannt, dass IT-Branche eine neue Branche und auch eine Kette ist, die die wirtschaftliche Entwicklung fördert. So ist ihre Position nicht zu ignorieren. Die Google Professional-Data-Engineer IT-Zertifizierung ist eine Methode für den Wettbewerb. Durch die Google Professional-Data-Engineer Zertifizierung werden Sie sich in allen Aspekten verbessern. Aber es ist nicht so einfach, die Prüfung zu bestehen. So empfehle ich Ihnen unsere originale Fragen. Wenn Sie die Schulungsressourcen wählen, ist ZertFragen die erste Wahl. Seine Erfolgsquote beträgt 100%. Und Sie können die Google Professional-Data-Engineer Prüfung sicher bestehen.

Die Google Professional-Data-Engineer-Prüfung ist eine Zertifizierungsprüfung, die von der Google Cloud Plattform angeboten wird. Sie richtet sich an Datenprofis, die ihre Expertise im Entwurf, Aufbau und Management von Datenverarbeitungssystemen auf der Google Cloud Plattform demonstrieren möchten. Die Prüfung testet das Wissen der Kandidaten über Datenengineering, Datenverarbeitung und Datenanalysetechniken unter Verwendung von Tools und Diensten der Google Cloud Plattform.

Die Google Professional-Data-Engineer-Prüfung ist eine der renommiertesten und begehrtesten Zertifizierungen im Bereich Data Engineering. Es wurde entwickelt, um das Wissen, die Fähigkeiten und das Fachwissen von Fachleuten beim Entwerfen, Aufbau und Wartung von Datenverarbeitungssystemen auf der Google Cloud -Plattform zu testen. Die Prüfung richtet sich an diejenigen, die Erfahrung im Datenentwicklung haben und ihre Fähigkeiten und ihr Wissen in diesem Bereich validieren möchten.

Um ein Google -zertifizierter professioneller Dateningenieur zu werden, muss ein Kandidat die Zertifizierungsprüfung bestehen, die

200 US -Dollar kostet. Die Prüfung ist in Englisch, Japanisch und Spanisch erhältlich und kann online oder in einem Testzentrum abgelegt werden. Die Prüfung ist zwei Jahre lang gültig, woraufhin ein Kandidat erneut wieder aufrechtzuerhalten muss, um seine Zertifizierung aufrechtzuerhalten.

## Google Certified Professional Data Engineer Exam Professional-Data-Engineer Prüfungsfragen mit Lösungen (Q85-Q90):

### 85. Frage

You need to create a SQL pipeline. The pipeline runs an aggregate SQL transformation on a BigQuery table every two hours and appends the result to another existing BigQuery table. You need to configure the pipeline to retry if errors occur. You want the pipeline to send an email notification after three consecutive failures.

What should you do?

- A. Create a BigQuery scheduled query to run the SQL transformation with schedule options that repeats every two hours, and enable emailnotifications.
- B. Use the BigQueryUpsertTableOperator in Cloud Composer, set the retry parameter to three, and set the email\_on\_failure parameter to true.
- C. Create a BigQuery scheduled query to run the SQL transformation with schedule options that repeats every two hours, and enable notification to Pub/Sub topic. Use Pub/Sub and Cloud Functions to send an email after three failed executions.
- D. Use the BigQueryInsertJobOperator in Cloud Composer, set the retry parameter to three, and set the email\_on\_failure parameter to true.

**Antwort: C**

Begründung:

To create a robust and resilient SQL pipeline in BigQuery that handles retries and failure notifications, consider the following:  
BigQuery Scheduled Queries: This feature allows you to schedule recurring queries in BigQuery. It is a straightforward way to run SQL transformations on a regular basis without requiring extensive setup.

Error Handling and Retries: While BigQuery Scheduled Queries can run at specified intervals, they don't natively support complex retry logic or failure notifications directly. This is where additional Google Cloud services like Pub/Sub and Cloud Functions come into play.

Pub/Sub for Notifications: By configuring a BigQuery scheduled query to publish messages to a Pub/Sub topic upon failure, you can create a decoupled and scalable notification system.

Cloud Functions: Cloud Functions can subscribe to the Pub/Sub topic and implement logic to count consecutive failures. After detecting three consecutive failures, the Cloud Function can then send an email notification using a service like SendGrid or Gmail API.

Implementation Steps:

Set up a BigQuery Scheduled Query:

Create a scheduled query in BigQuery to run your SQL transformation every two hours.

Configure the scheduled query to publish a notification to a Pub/Sub topic in case of a failure.

Create a Pub/Sub Topic:

Create a Pub/Sub topic that will receive messages from the scheduled query.

Develop a Cloud Function:

Write a Cloud Function that subscribes to the Pub/Sub topic.

Implement logic in the Cloud Function to track failure messages. If three consecutive failure messages are detected, the function sends an email notification.

Reference Links:

BigQuery Scheduled Queries

Pub/Sub Documentation

Cloud Functions Documentation

SendGrid Email API

Gmail API

### 86. Frage

Your company is currently setting up data pipelines for their campaign. For all the Google Cloud Pub/Sub streaming data, one of the important business requirements is to be able to periodically identify the inputs and their timings during their campaign. Engineers have decided to use windowing and transformation in Google Cloud Dataflow for this purpose. However, when testing this feature, they find that the Cloud Dataflow job fails for all streaming inserts. What is the most likely cause of this problem?

- A. They have not set the triggers to accommodate the data coming in late, which causes the job to fail
- **B. They have not applied a global windowing function, which causes the job to fail when the pipeline is created**
- C. They have not applied a non-global windowing function, which causes the job to fail when the pipeline is created
- D. They have not assigned the timestamp, which causes the job to fail

**Antwort: B**

### 87. Frage

You are designing a data mesh on Google Cloud with multiple distinct data engineering teams building data products. The typical data curation design pattern consists of landing files in Cloud Storage, transforming raw data in Cloud Storage and BigQuery datasets, and storing the final curated data product in BigQuery datasets. You need to configure Dataplex to ensure that each team can access only the assets needed to build their data products. You also need to ensure that teams can easily share the curated data product. What should you do?

- **A. 1 Create a Dataplex virtual lake for each data product, and create multiple zones for landing, raw, and curated data. 2. Provide the data engineering teams with full access to the virtual lake assigned to their data product.**
- B. 1 Create a single Dataplex virtual lake and create a single zone to contain landing, raw, and curated data. 2 Provide each data engineering team access to the virtual lake.
- C. 1 Create a single Dataplex virtual lake and create a single zone to contain landing, raw, and curated data. 2 Build separate assets for each data product within the zone. 3. Assign permissions to the data engineering teams at the zone level.
- D. 1 Create a Dataplex virtual lake for each data product, and create a single zone to contain landing, raw, and curated data. 2. Provide the data engineering teams with full access to the virtual lake assigned to their data product.

**Antwort: A**

Begründung:

This option is the best way to configure Dataplex for a data mesh architecture, as it allows each data engineering team to have full ownership and control over their data products, while also enabling easy discovery and sharing of the curated data across the organization<sup>1,2</sup>. By creating a Dataplex virtual lake for each data product, you can isolate the data assets and resources for each domain, and avoid conflicts and dependencies between different teams<sup>3</sup>. By creating multiple zones for landing, raw, and curated data, you can enforce different security and governance policies for each stage of the data curation process, and ensure that only authorized users can access the data assets<sup>4,5</sup>. By providing the data engineering teams with full access to the virtual lake assigned to their data product, you can empower them to manage and monitor their data products, and leverage the Dataplex features such as tagging, quality, and lineage.

Option A is not suitable, as it creates a single point of failure and a bottleneck for the data mesh, and does not allow for fine-grained access control and governance for different data products<sup>2</sup>. Option B is also not suitable, as it does not isolate the data assets and resources for each data product, and assigns permissions at the zone level, which may not reflect the different roles and responsibilities of the data engineering teams<sup>3,4</sup>. Option C is better than option A and B, but it does not create multiple zones for landing, raw, and curated data, which may compromise the security and quality of the data products<sup>5</sup>. References:

\* 1: Building a data mesh on Google Cloud using BigQuery and Dataplex | Google Cloud Blog

\* 2: Data Mesh - 7 Effective Practices to Get Started - Confluent

\* 3: Best practices | Dataplex | Google Cloud

\* 4: Secure your lake | Dataplex | Google Cloud

\* 5: Zones | Dataplex | Google Cloud

\* [6]: Managing a Data Mesh with Dataplex - ROI Training

### 88. Frage

Government regulations in your industry mandate that you have to maintain an auditable record of access to certain types of data. Assuming that all expiring logs will be archived correctly, where should you store data that is subject to that mandate?

- A. In a bucket on Cloud Storage that is accessible only by an AppEngine service that collects user information and logs the access before providing a link to the bucket.
- B. In Cloud SQL, with separate database user names to each user. The Cloud SQL Admin activity logs will be used to provide the auditability.
- **C. In a BigQuery dataset that is viewable only by authorized personnel, with the Data Access log used to provide the auditability.**
- D. Encrypted on Cloud Storage with user-supplied encryption keys. A separate decryption key will be

given to each authorized user.

**Antwort: C**

## 89. Frage

Case Study 2 - MJTelco

Company Overview

MJTelco is a startup that plans to build networks in rapidly growing, underserved markets around the world.

The company has patents for innovative optical communications hardware. Based on these patents, they can create many reliable, high-speed backbone links with inexpensive hardware.

Company Background

Founded by experienced telecom executives, MJTelco uses technologies originally developed to overcome communications challenges in space. Fundamental to their operation, they need to create a distributed data infrastructure that drives real-time analysis and incorporates machine learning to continuously optimize their topologies. Because their hardware is inexpensive, they plan to overdeploy the network allowing them to account for the impact of dynamic regional politics on location availability and cost. Their management and operations teams are situated all around the globe creating many-to-many relationship between data consumers and provides in their system. After careful consideration, they decided public cloud is the perfect environment to support their needs.

Solution Concept

MJTelco is running a successful proof-of-concept (PoC) project in its labs. They have two primary needs:

- \* Scale and harden their PoC to support significantly more data flows generated when they ramp to more than 50,000 installations.
- \* Refine their machine-learning cycles to verify and improve the dynamic models they use to control topology definition.

MJTelco will also use three separate operating environments - development/test, staging, and production - to meet the needs of running experiments, deploying new features, and serving production customers.

Business Requirements

- \* Scale up their production environment with minimal cost, instantiating resources when and where needed in an unpredictable, distributed telecom user community.
- \* Ensure security of their proprietary data to protect their leading-edge machine learning and analysis.
- \* Provide reliable and timely access to data for analysis from distributed research workers
- \* Maintain isolated environments that support rapid iteration of their machine-learning models without affecting their customers.

Technical Requirements

- \* Ensure secure and efficient transport and storage of telemetry data
- \* Rapidly scale instances to support between 10,000 and 100,000 data providers with multiple flows each.
- \* Allow analysis and presentation against data tables tracking up to 2 years of data storing approximately 100m records/day
- \* Support rapid iteration of monitoring infrastructure focused on awareness of data pipeline problems both in telemetry flows and in production learning cycles.

CEO Statement

Our business model relies on our patents, analytics and dynamic machine learning. Our inexpensive hardware is organized to be highly reliable, which gives us cost advantages. We need to quickly stabilize our large distributed data pipelines to meet our reliability and capacity commitments.

CTO Statement

Our public cloud services must operate as advertised. We need resources that scale and keep our data secure. We also need environments in which our data scientists can carefully study and quickly adapt our models. Because we rely on automation to process our data, we also need our development and test environments to work as we iterate.

CFO Statement

The project is too large for us to maintain the hardware and software required for the data and analysis.

Also, we cannot afford to staff an operations team to monitor so many data feeds, so we will rely on automation and infrastructure. Google Cloud's machine learning will allow our quantitative researchers to work on our high-value problems instead of problems with our data pipelines.

Given the record streams MJTelco is interested in ingesting per day, they are concerned about the cost of Google BigQuery increasing. MJTelco asks you to provide a design solution. They require a single large data table called `tracking_table`. Additionally, they want to minimize the cost of daily queries while performing fine-grained analysis of each day's events. They also want to use streaming ingestion. What should you do?

- A. Create a table called `tracking_table` with a `TIMESTAMP` column to represent the day.
- **B. Create a partitioned table called `tracking_table` and include a `TIMESTAMP` column.**
- C. Create sharded tables for each day following the pattern `tracking_table_YYYYMMDD`.
- D. Create a table called `tracking_table` and include a `DATE` column.

**Antwort: B**

## 90. Frage

.....

Die Prüfungsunterlagen zur Google Professional-Data-Engineer Zertifizierungsprüfung werden nach dem Lehrkompendium und den echten Prüfungen bearbeitet. Wir aktualisieren auch ständig unsere Schulungsunterlagen, so dass Sie in erster Zeit die neuesten und besten Informationen bekommen. Wenn Sie unsere Schulungsunterlagen zur Google Professional-Data-Engineer Zertifizierungsprüfung kaufen, können Sie einen einjährigen kostenlosen Update-Service bekommen. Sie können jederzeit Abonnementzeit verlängern, so dass Sie mehr Zeit haben, sich auf die Google Professional-Data-Engineer Prüfung vorzubereiten.

**Professional-Data-Engineer Schulungsangebot:** [https://www.zertfragen.com/Professional-Data-Engineer\\_pruefung.html](https://www.zertfragen.com/Professional-Data-Engineer_pruefung.html)

- Sie können so einfach wie möglich - Professional-Data-Engineer bestehen! ☐ Suchen Sie auf ☐ [www.zertfragen.com](http://www.zertfragen.com) ☐ nach kostenlosem Download von ☐ Professional-Data-Engineer ☐ iProfessional-Data-Engineer Vorbereitungsfragen
- Professional-Data-Engineer Zertifizierungsprüfung ☐ Professional-Data-Engineer Tests ☐ Professional-Data-Engineer Vorbereitungsfragen ☐ Sie müssen nur zu ☐ [www.itzert.com](http://www.itzert.com) ☐ gehen um nach kostenloser Download von ☐ Professional-Data-Engineer ☐ zu suchen ☐ Professional-Data-Engineer Antworten
- Professional-Data-Engineer Schulungsangebot - Professional-Data-Engineer Simulationsfragen - Professional-Data-Engineer kostenlos downloaden ☐ Sie müssen nur zu ☐ [www.zertsoft.com](http://www.zertsoft.com) ☐ gehen um nach kostenloser Download von ☐ Professional-Data-Engineer ☐ zu suchen ☐ Professional-Data-Engineer Fragen&Antworten
- Professional-Data-Engineer Fragen Und Antworten ☐ Professional-Data-Engineer Übungsmaterialien ☐ Professional-Data-Engineer Antworten ☐ ☐ [www.itzert.com](http://www.itzert.com) ☐ ist die beste Webseite um den kostenlosen Download von ☐ Professional-Data-Engineer ☐ zu erhalten ☐ Professional-Data-Engineer Übungsmaterialien
- Professional-Data-Engineer Dumps und Test Überprüfungen sind die beste Wahl für Ihre Google Professional-Data-Engineer Testvorbereitung ☐ Öffnen Sie die Website ☐ [www.zertpruefung.de](http://www.zertpruefung.de) ☐ Suchen Sie ☒ Professional-Data-Engineer ☐ ☒ Kostenloser Download ☐ Professional-Data-Engineer Schulungsangebot
- Professional-Data-Engineer Torrent Anleitung - Professional-Data-Engineer Studienführer - Professional-Data-Engineer wirkliche Prüfung ☐ Öffnen Sie die Webseite ☒ [www.itzert.com](http://www.itzert.com) ☒ und suchen Sie nach kostenloser Download von ☒ Professional-Data-Engineer ☐ ☒ Professional-Data-Engineer Übungsmaterialien
- Professional-Data-Engineer Fragen Und Antworten ☐ Professional-Data-Engineer Fragen&Antworten ☐ Professional-Data-Engineer Deutsche ☐ Suchen Sie einfach auf ☐ [www.zertpruefung.ch](http://www.zertpruefung.ch) ☐ nach kostenloser Download von ☐ Professional-Data-Engineer ☐ Professional-Data-Engineer Fragen&Antworten
- Sie können so einfach wie möglich - Professional-Data-Engineer bestehen! ☐ ☐ [www.itzert.com](http://www.itzert.com) ☐ ist die beste Webseite um den kostenlosen Download von ☐ Professional-Data-Engineer ☐ zu erhalten ☐ Professional-Data-Engineer Tests
- Professional-Data-Engineer Prüfungsfragen Prüfungsvorbereitungen 2025: Google Certified Professional Data Engineer Exam - Zertifizierungsprüfung Google Professional-Data-Engineer in Deutsch Englisch pdf downloaden ☐ Geben Sie ☐ [www.zertfragen.com](http://www.zertfragen.com) ☐ ein und suchen Sie nach kostenloser Download von ☐ Professional-Data-Engineer ☐ ☐ ☐ Professional-Data-Engineer Lernhilfe
- Professional-Data-Engineer Schulungsangebot - Professional-Data-Engineer Simulationsfragen - Professional-Data-Engineer kostenlos downloaden ☐ Öffnen Sie die Website ☐ [www.itzert.com](http://www.itzert.com) ☐ Suchen Sie ☐ Professional-Data-Engineer ☐ Kostenloser Download ☐ Professional-Data-Engineer Deutsch Prüfung
- Professional-Data-Engineer Deutsche ☐ Professional-Data-Engineer Originale Fragen ☐ Professional-Data-Engineer Schulungsangebot ☐ Suchen Sie jetzt auf ☐ [www.deutschpruefung.com](http://www.deutschpruefung.com) ☐ nach ☐ Professional-Data-Engineer ☐ und laden Sie es kostenlos herunter ☐ Professional-Data-Engineer Simulationsfragen
- [www.stes.tyc.edu.tw](http://www.stes.tyc.edu.tw), [www.stes.tyc.edu.tw](http://www.stes.tyc.edu.tw), [infodots.in](http://infodots.in), [www.stes.tyc.edu.tw](http://www.stes.tyc.edu.tw), [www.stes.tyc.edu.tw](http://www.stes.tyc.edu.tw), [karenobrien1.blogspot.com](http://karenobrien1.blogspot.com), [joshha964.webdesign96.com](http://joshha964.webdesign96.com), [lms.ait.edu.za](http://lms.ait.edu.za), [hi-wot.com](http://hi-wot.com), [shortcourses.russellcollege.edu.au](http://shortcourses.russellcollege.edu.au), Disposable vapes

**BONUS!!!** Laden Sie die vollständige Version der ZertFragen Professional-Data-Engineer Prüfungsfragen kostenlos herunter:  
[https://drive.google.com/open?id=1wmHqxRJhud83KTHZ1i\\_ELkfPYL7zUHD5](https://drive.google.com/open?id=1wmHqxRJhud83KTHZ1i_ELkfPYL7zUHD5)