



VirtuLearn
Trusted, Accurate & Reliable
vp

TRUSTED, ACCURATE AND RELIABLE

The most comprehensive IT certification
preparation materials in the industry!



Expert-Verified



Exam-Ready



Regularly Updated

All rights reserved. No part of this document may be reproduced, distributed, or transmitted in any form or by any means, including photocopying, recording, or other electronic or mechanical methods, without the prior written permission of the publisher, except in the case of brief quotations embodied in critical reviews and certain other non-commercial uses permitted by copyright law. Unauthorized copying, reselling, or distribution of this document is strictly prohibited and may result in legal action.

Snowflake

GES-C01

SnowPro Specialty Gen AI
Certification Exam

QUESTION: 1

What privileges must be granted to a user with the FNR_DEVELOP_ROLE so the user can use Document AI on an existing pipeline? (Choose three.)

- A. GRANT DATABASE ROLE SNOWFLAKE.ML_USER TO ROLE FNR_DEVELOP_ROLE;
- B. GRANT DATABASE ROLE SNOWFLAKE.DOCUMENT_INTELLIGENCE_CREATOR TO ROLE FNR_DEVELOP_ROLE;
- C. GRANT ROLE SNOWFLAKE.DOCUMENT_INTELLIGENCE_CREATOR TO ROLE FNR_DEVELOP_ROLE;
- D. GRANT CREATE SNOWFLAKE.ML.DOCUMENT_INTELLIGENCE ON SCHEMA <database>.<schema> TO ROLE FNR_DEVELOP_ROLE;
- E. GRANT CREATE SNOWFLAKE.MODELS.DOCUMENT_INTELLIGENCE ON SCHEMA <database>.<schema>
- F. GRANT CREATE MODEL ON SCHEMA <database>.<schema> TO ROLE FNR_DEVELOP_ROLE;

Answer(s): B, C, F

Explanation:**Justification (structured, bullet-point style)**

B – GRANT DATABASE ROLE SNOWFLAKE Grants the built-in database-level role that includes the USAGE privilege on the target database.

This is the minimum privilege needed for a user to see and reference objects (tables, schemas, stages) that belong to the database housing the Document AI pipeline. Without USAGE on the database, the pipeline's metadata cannot be located, so the user cannot even start a Document AI execution.

C – GRANT ROLE SNOWFLAKE Allows the user to create transient session-level objects (e.g., temporary tables, internal stages) within a schema that the pipeline may use. Document AI pipelines often materialize intermediate results in temporary schema objects. The GRANT ROLE

SNOWFLAKE (or equivalent CREATE on the schema) provides the CREATE privilege required to create those objects. Without it, the pipeline would fail when trying to write temporary data.

F – MODELS.DOCUMENT_INTELLIGENCE ON SCHEMA <database>.<schema> Explicitly grants the model-level privilege that enables invocation of the Document AI model service for pipelines referencing a model stored in that schema. The pipeline must be authorized to call the underlying ML model. Granting MODELS.DOCUMENT_INTELLIGENCE on the schema gives the role the USAGE privilege on the model object,

which is required for the SELECT -style calls that Document AI performs. This is the only privilege that directly authorizes interaction with the Document AI service.

Why the other options are unsuitable

A – GRANT DATABASE ROLE SNOWFLAK (as listed) – The syntax is incomplete; the correct grant is GRANT

DATABASE ROLE SNOWFLAKE **or** GRANT ROLE SNOWFLAKE TO USER As written it does not actually grant any privilege, so it cannot satisfy the requirement.

D – GRANT CREATE SNOWFLAK – CREATE SNOWFLAK is not a valid Snowflake privilege; the only supported privilege is CREATE on a **schema** (or CREATE MODEL on a **model**). This statement therefore represents an invalid privilege and cannot be used.

E – DOCUMENT_INTELLIGENCE_CREATOR / ML_USER – These are higher-level roles intended for developers building new Document AI models. Granting them would give the user far more rights (e.g., CREATE MODEL, CREATE PIPELINE, etc.) than are needed for merely consuming an existing pipeline. They over-privilege the user and are not part of the minimal privilege set.

Consequently, the minimal and correct combination is **B, C, and F** – the database-level access, the ability to create temporary schema objects, and the explicit model-usage grant.

Reference:

Snowflake Documentation – Granting Database and Schema Privileges:

<https://docs.snowflake.com/en/sql-reference/sql/grant-privileges> Snowflake Documentation –

Document AI – Model Privileges: <https://docs.snowflake.com/en/user-guide/snowpark-container-services/document-ai#model-privileges>

QUESTION: 2

When using the SNOWFLAKE.CORTEX.FINETUNE function, which column names must be included in the SQL statement in the training_data_query parameter? (Choose two.)

- A. BASE_MODEL
- B. COMPLETION
- C. INFERENCE
- D. PROMPT
- E. TRAINING

Answer(s): B, D

Explanation:

Why option BD is correct

The SNOWFLAKE.CORTEX.FINETUNE function expects the training_data_query argument to be a SELECT

statement that returns **exactly two columns** that the model uses as input and output signals.

The valid column identifiers required by the function are PROMPT (the instruction or user input) and

COMPLETION (the desired response). These are referenced internally as the “B” and “D” labels in the exam’s answer set.

In the SQL you must explicitly include these column names in the SELECT list, e.g.:

```
SELECT prompt AS prompt, -- column B completion AS completion -- column D
FROM raw_training_data;
```

If either of these columns is missing or has a different alias, the function throws a validation error because the model cannot map the training rows to the expected input-output format.

Why the other option is not suitable