

CERT-PASS

Snowflake SnowPro Advanced: Data Engineer

Free Practice Questions Preview

Here are 35 sample questions to help you get started. Unlock the full exam to access all 1050+ questions with detailed explanations.

Question 1 : Domain 1.0: Data Movement

[Easy] The SaaS vendor is preparing a production rollout for retryable extraction from an external API. The team wants to meet the SLA without overengineering. The engineer must choose the option whose tradeoff matches the requirement. Which change should the engineer recommend?

- A. Store API credentials in a table and expose the table to all pipeline roles
- B. Land API extracts in a controlled stage or landing table, then use an idempotent load and transformation path**
- C. Call the external API from a row access policy during every analytical query
- D. Use a secure share as the API client and expect it to persist responses

Answer: B

Correct answer: B. A controlled landing pattern makes extraction retryable, observable, and separable from transformation. Why the strongest distractor fails: C. Call the external API from a row access policy during every analytical query — Invoking APIs during analytical queries couples availability and latency to user workloads and is not a resilient ingestion design.

Question 2 : Domain 1.0: Data Movement

[Easy] An architecture board for the travel marketplace evaluates retryable extraction from an external API. The proposed fix must avoid copying data unnecessarily. A pull request uses a plausible but incorrect alternative. Which option best satisfies the requirement?

- A. Land API extracts in a controlled stage or landing table, then use an idempotent load and transformation path**
- B. Call the external API from a row access policy during every analytical query
- C. Use a secure share as the API client and expect it to persist responses
- D. Store API credentials in a table and expose the table to all pipeline roles

Answer: A

Correct answer: A. A controlled landing pattern makes extraction retryable, observable, and separable from transformation. Why the strongest distractor fails: B. Call the external API from a row access policy during every analytical query — Invoking APIs during analytical queries couples availability and latency to user workloads and is not a resilient ingestion design.

Question 3 : Domain 1.0: Data Movement

[Easy] The supply-chain program is preparing a production rollout for retryable extraction from an external API. The team wants to keep the solution idempotent. The team is ready to implement one targeted change. Which approach should be selected?

- A. Call the external API from a row access policy during every analytical query
- B. Use a secure share as the API client and expect it to persist responses
- C. Store API credentials in a table and expose the table to all pipeline roles
- D. Land API extracts in a controlled stage or landing table, then use an idempotent load and transformation path**

Answer: D

Correct answer: D. A controlled landing pattern makes extraction retryable, observable, and separable from transformation. Why the strongest distractor fails: A. Call the external API from a row access policy during every analytical query — Invoking APIs during analytical queries couples availability and latency to user workloads and is not a resilient ingestion design.

Question 4 : Domain 1.0: Data Movement

[Easy] An architecture board for the financial reporting team evaluates retryable extraction from an external API. The proposed fix must meet the SLA without overengineering. Which statement identifies the correct Snowflake capability?

- A. Use a secure share as the API client and expect it to persist responses
- B. Store API credentials in a table and expose the table to all pipeline roles
- C. Land API extracts in a controlled stage or landing table, then use an idempotent load and transformation path**
- D. Call the external API from a row access policy during every analytical query

Answer: C

Correct answer: C. A controlled landing pattern makes extraction retryable, observable, and separable from transformation. Why the strongest distractor fails: D. Call the external API from a row access policy during every analytical query — Invoking APIs during analytical queries couples availability and latency to user workloads and is not a resilient ingestion design.

Question 5 : Domain 1.0: Data Movement

[Easy] During a production incident at the public-sector reporting unit, engineers examine a staged-file bulk load that may be retried after orchestration failures. They need to retain useful diagnostic evidence. The first proposed remediation targets the wrong Snowflake feature. What is the strongest next action?

- A. Run INSERT SELECT from the stage on every schedule without a file manifest
- B. Use COPY INTO <location> to unload the files and reload the exported copies
- C. Create a secure share for the stage and let consumers insert the records
- D. COPY INTO <table> from the stage and rely on table load metadata for previously loaded files**

Answer: D

Correct answer: D. COPY INTO <table> records file-loading metadata for a target table, which is the native control for avoiding accidental reprocessing of staged files. Why the strongest distractor fails: B. Use COPY INTO <location> to unload the files and reload the exported copies — INSERT SELECT does not provide the same COPY load-history behavior and can reload files unless the pipeline adds its own controls.

Question 6 : Domain 1.0: Data Movement

[Easy] A code review for the healthcare analytics team finds a weak approach in a staged-file bulk load that may be retried after orchestration failures. The corrected approach should use native Snowflake capabilities. The engineer must choose the option whose tradeoff matches the requirement. What should the team implement?

- A. Use COPY INTO <location> to unload the files and reload the exported copies
- B. Create a secure share for the stage and let consumers insert the records
- C. COPY INTO <table> from the stage and rely on table load metadata for previously loaded files**
- D. Run INSERT SELECT from the stage on every schedule without a file manifest

Answer: C

Correct answer: C. COPY INTO <table> records file-loading metadata for a target table, which is the native control for avoiding accidental reprocessing of staged files. Why the strongest distractor fails: A. Use COPY INTO <location> to unload the files and reload the exported copies — INSERT SELECT does not provide the same COPY load-history behavior and can reload files unless the pipeline adds its own controls.

Question 7 : Domain 1.0: Data Movement

[Easy] During a production incident at the customer-support analytics team, engineers examine a staged-file bulk load that may be retried after orchestration failures. They need to preserve least-privilege access. A pull request uses a plausible but incorrect alternative. Which change should the engineer recommend?

- A. Create a secure share for the stage and let consumers insert the records
- B. COPY INTO <table> from the stage and rely on table load metadata for previously loaded files**
- C. Run INSERT SELECT from the stage on every schedule without a file manifest
- D. Use COPY INTO <location> to unload the files and reload the exported copies

Answer: B

Correct answer: B. COPY INTO <table> records file-loading metadata for a target table, which is the native control for avoiding accidental reprocessing of staged files. Why the strongest distractor fails: D. Use COPY INTO <location> to unload the files and reload the exported copies — INSERT SELECT does not provide the same COPY load-history behavior and can reload files unless the pipeline adds its own controls.

Question 8 : Domain 1.0: Data Movement

[Easy] A code review for the insurance data platform finds a weak approach in a staged-file bulk load that may be retried after orchestration failures. The corrected approach should retain useful diagnostic evidence. The team is ready to implement one targeted change. Which option best satisfies the requirement?

- A. COPY INTO <table> from the stage and rely on table load metadata for previously loaded files**
- B. Run INSERT SELECT from the stage on every schedule without a file manifest
- C. Use COPY INTO <location> to unload the files and reload the exported copies
- D. Create a secure share for the stage and let consumers insert the records

Answer: A

Correct answer: A. COPY INTO <table> records file-loading metadata for a target table, which is the native control for avoiding accidental reprocessing of staged files. Why the strongest distractor fails: C. Use COPY INTO <location> to unload the files and reload the exported copies — INSERT SELECT does not provide the same COPY load-history behavior and can reload files unless the pipeline adds its own controls.

Question 9 : Domain 1.0: Data Movement

[Easy] A senior data engineer at the e-commerce marketplace is redesigning source files whose stable headers arrive in a different order. The design should preserve predictable behavior during retries. A pull request uses a plausible but incorrect alternative. Which correction addresses the requirement most directly?

- A. Use MATCH_BY_COLUMN_NAME when loading files whose column order can vary but column names are stable
- B. Set FORCE = TRUE so column order is ignored
- C. Increase the warehouse size so COPY can infer column names from position
- D. Use PURGE = TRUE to remap the source fields

Answer: A

Correct answer: A. MATCH_BY_COLUMN_NAME aligns file fields to target columns by name rather than relying on ordinal position. Why the strongest distractor fails: B. Set FORCE = TRUE so column order is ignored — FORCE controls whether previously loaded files may be reloaded; it does not change column mapping.

Question 10 : Domain 1.0: Data Movement

[Easy] After a failed test run, the logistics network revisits source files whose stable headers arrive in a different order. The engineer must support auditability. The team is ready to implement one targeted change. Which design is the best fit?

- A. Set FORCE = TRUE so column order is ignored
- B. Increase the warehouse size so COPY can infer column names from position
- C. Use PURGE = TRUE to remap the source fields
- D. Use MATCH_BY_COLUMN_NAME when loading files whose column order can vary but column names are stable

Answer: D

Correct answer: D. MATCH_BY_COLUMN_NAME aligns file fields to target columns by name rather than relying on ordinal position. Why the strongest distractor fails: A. Set FORCE = TRUE so column order is ignored — FORCE controls whether previously loaded files may be reloaded; it does not change column mapping.

Question 11 : Domain 1.0: Data Movement

[Easy] A senior data engineer at the manufacturing group is redesigning source files whose stable headers arrive in a different order. The design should minimize operational overhead. Which statement identifies the correct Snowflake capability?

- A. Increase the warehouse size so COPY can infer column names from position

B. Use PURGE = TRUE to remap the source fields

C. Use MATCH_BY_COLUMN_NAME when loading files whose column order can vary but column names are stable

D. Set FORCE = TRUE so column order is ignored

Answer: C

Correct answer: C. MATCH_BY_COLUMN_NAME aligns file fields to target columns by name rather than relying on ordinal position. Why the strongest distractor fails: D. Set FORCE = TRUE so column order is ignored — FORCE controls whether previously loaded files may be reloaded; it does not change column mapping.

Question 12 : Domain 1.0: Data Movement

[Easy] After a failed test run, the retail platform revisits source files whose stable headers arrive in a different order. The engineer must preserve predictable behavior during retries. The existing design works in a small test but is not acceptable for production. What should the team implement?

A. Use PURGE = TRUE to remap the source fields

B. Use MATCH_BY_COLUMN_NAME when loading files whose column order can vary but column names are stable

C. Set FORCE = TRUE so column order is ignored

D. Increase the warehouse size so COPY can infer column names from position

Answer: B

Correct answer: B. MATCH_BY_COLUMN_NAME aligns file fields to target columns by name rather than relying on ordinal position. Why the strongest distractor fails: C. Set FORCE = TRUE so column order is ignored — FORCE controls whether previously loaded files may be reloaded; it does not change column mapping.

Question 13 : Domain 1.0: Data Movement

[Easy] The supply-chain program is preparing a production rollout for a writable continuity copy across Snowflake accounts. The team wants to keep the solution idempotent. Which statement identifies the correct Snowflake capability?

A. Configure replication or failover groups for the required databases and account objects across accounts

B. Use a secure share as a disaster-recovery copy that can be promoted for writes

C. Use Search Optimization Service to synchronize the secondary account

D. Clone the production database into a different cloud region with a single zero-copy clone statement

Answer: A

Correct answer: A. Replication and failover groups are designed to synchronize data and supported account objects across regions and clouds. Why the strongest distractor fails: B. Use a secure share as a disaster-recovery copy that can be promoted for writes — A secure share provides governed read access; it is not a writable failover copy.

Question 14 : Domain 1.0: Data Movement

[Easy] The retail platform is reviewing a writable continuity copy across Snowflake accounts. The priority is to preserve predictable behavior during retries. The team is ready to implement one targeted change. Which implementation is most appropriate?

- A. Use Search Optimization Service to synchronize the secondary account
- B. Use a secure share as a disaster-recovery copy that can be promoted for writes
- C. Configure replication or failover groups for the required databases and account objects across accounts**
- D. Clone the production database into a different cloud region with a single zero-copy clone statement

Answer: C

Correct answer: C. Replication and failover groups are designed to synchronize data and supported account objects across regions and clouds. Why the strongest distractor fails: B. Use a secure share as a disaster-recovery copy that can be promoted for writes — A secure share provides governed read access; it is not a writable failover copy.

Question 15 : Domain 1.0: Data Movement

[Easy] For the e-commerce marketplace, a new requirement affects a writable continuity copy across Snowflake accounts. The implementation must support auditability. Which statement identifies the correct Snowflake capability?

- A. Clone the production database into a different cloud region with a single zero-copy clone statement
- B. Use Search Optimization Service to synchronize the secondary account
- C. Use a secure share as a disaster-recovery copy that can be promoted for writes
- D. Configure replication or failover groups for the required databases and account objects across accounts**

Answer: D

Correct answer: D. Replication and failover groups are designed to synchronize data and supported account objects across regions and clouds. Why the strongest distractor fails: C. Use a secure share as a disaster-recovery copy that can be promoted for writes — A secure share provides governed read access; it is not a writable failover copy.

Question 16 : Domain 1.0: Data Movement

[Easy] The logistics network is reviewing a writable continuity copy across Snowflake accounts. The priority is to minimize operational overhead. The engineer should not confuse storage, compute, orchestration, and governance responsibilities. Which design is the best fit?

- A. Configure replication or failover groups for the required databases and account objects across accounts**
- B. Clone the production database into a different cloud region with a single zero-copy clone statement
- C. Use Search Optimization Service to synchronize the secondary account
- D. Use a secure share as a disaster-recovery copy that can be promoted for writes

Answer: A

Correct answer: A. Replication and failover groups are designed to synchronize data and supported account objects across regions and clouds. Why the strongest distractor fails: D. Use a secure share as a disaster-recovery copy that can be promoted for writes — A secure share provides governed read access; it is not a writable failover copy.

Question 17 : Domain 1.0: Data Movement

[Easy] For the telecommunications provider, a new requirement affects a workflow that needs queryable metadata for staged files. The implementation must keep recovery behavior explicit. Which statement identifies the correct Snowflake capability?

- A. Use a masking policy to expose file names
- B. Enable a directory table on the stage when the pipeline needs metadata about staged files**
- C. Enable Time Travel on the stage to query its file list
- D. Create a materialized view over LIST output

Answer: B

Correct answer: B. A directory table exposes metadata about files in a stage for querying and processing workflows. Why the strongest distractor fails: C. Enable Time Travel on the stage to query its file list — Time Travel applies to supported Snowflake objects and is not a file catalog for a stage.

Question 18 : Domain 1.0: Data Movement

[Easy] The ad-tech platform is reviewing a workflow that needs queryable metadata for staged files. The priority is to avoid unnecessary credit consumption. A reviewer rejects any answer that solves a different layer of the architecture. Which implementation is most appropriate?

- A. Enable a directory table on the stage when the pipeline needs metadata about staged files**
- B. Enable Time Travel on the stage to query its file list
- C. Create a materialized view over LIST output
- D. Use a masking policy to expose file names

Answer: A

Correct answer: A. A directory table exposes metadata about files in a stage for querying and processing workflows. Why the strongest distractor fails: B. Enable Time Travel on the stage to query its file list — Time Travel applies to supported Snowflake objects and is not a file catalog for a stage.

Question 19 : Domain 1.0: Data Movement

[Easy] For the online learning provider, a new requirement affects a workflow that needs queryable metadata for staged files. The implementation must separate compute concerns from storage concerns. Select the option containing the best primary action and the most relevant validation step.

- A. Enable Time Travel on the stage to query its file list; then test the design with inserts, updates, deletes, and replay where applicable
- B. Create a materialized view over LIST output; then verify the persistence and recovery tradeoff matches the data classification
- C. Use a masking policy to expose file names; then compare query profile metrics and credit usage before and after the change
- D. Enable a directory table on the stage when the pipeline needs metadata about staged files; then test the behavior with duplicate delivery and malformed input**

Answer: D

Correct answer: D. A directory table exposes metadata about files in a stage for querying and processing workflows.

Why the strongest distractor fails: A. Enable Time Travel on the stage to query its file list; then test the design with inserts, updates, deletes, and replay where applicable — Time Travel applies to supported Snowflake objects and is not a file catalog for a stage.

Question 20 : Domain 1.0: Data Movement

[Easy] The cybersecurity operations team is reviewing a workflow that needs queryable metadata for staged files. The priority is to keep recovery behavior explicit. The first proposed remediation targets the wrong Snowflake feature. Which design is the best fit?

- A. Create a materialized view over LIST output
- B. Use a masking policy to expose file names
- C. Enable a directory table on the stage when the pipeline needs metadata about staged files**
- D. Enable Time Travel on the stage to query its file list

Answer: C

Correct answer: C. A directory table exposes metadata about files in a stage for querying and processing workflows.

Why the strongest distractor fails: D. Enable Time Travel on the stage to query its file list — Time Travel applies to supported Snowflake objects and is not a file catalog for a stage.

Question 21 : Domain 1.0: Data Movement

[Easy] For the e-commerce marketplace, a new requirement affects an external-table metadata refresh design. The implementation must support auditability. The team needs a maintainable solution that can be explained during an audit. Which approach should be selected?

- A. Set PURGE = TRUE to convert invalid records to NULL automatically
- B. Set FORCE = TRUE to skip malformed records but keep valid rows
- C. Set ON_ERROR according to the batch SLA, such as CONTINUE for quarantining bad rows while loading valid rows**
- D. Set MATCH_BY_COLUMN_NAME = CASE_INSENSITIVE to ignore all data quality errors

Answer: C

Correct answer: C. ON_ERROR controls how COPY behaves when parsing or conversion errors occur and should match the pipeline SLA. Why the strongest distractor fails: B. Set FORCE = TRUE to skip malformed records but keep valid rows — FORCE affects reloading files already recorded in load metadata; it does not control row-level parsing errors.

Question 22 : Domain 1.0: Data Movement

[Easy] The logistics network is reviewing a consumer onboarding request. The priority is to minimize operational overhead. Select the option containing the best primary action and the most relevant validation step.

- A. Set MATCH_BY_COLUMN_NAME = CASE_INSENSITIVE to ignore all data quality errors; then monitor workload history rather than relying on one isolated query
- B. Set PURGE = TRUE to convert invalid records to NULL automatically; then re-run representative workloads and examine queueing, scan, and spill metrics
- C. Set FORCE = TRUE to skip malformed records but keep valid rows; then test the design with inserts, updates, deletes, and replay where applicable
- D. Set ON_ERROR according to the batch SLA, such as CONTINUE for quarantining bad rows while loading valid rows; then review ingestion history and retry behavior**

Answer: D

Correct answer: D. ON_ERROR controls how COPY behaves when parsing or conversion errors occur and should match the pipeline SLA. Why the strongest distractor fails: C. Set FORCE = TRUE to skip malformed records but keep valid rows; then test the design with inserts, updates, deletes, and replay where applicable — FORCE affects reloading files already recorded in load metadata; it does not control row-level parsing errors.

Question 23 : Domain 1.0: Data Movement

[Easy] For the manufacturing group, a new requirement affects a cloud-storage integration review. The implementation must preserve predictable behavior during retries. The first proposed remediation targets the wrong Snowflake feature. Which correction addresses the requirement most directly?

- A. Set ON_ERROR according to the batch SLA, such as CONTINUE for quarantining bad rows while loading valid rows**
- B. Set MATCH_BY_COLUMN_NAME = CASE_INSENSITIVE to ignore all data quality errors
- C. Set PURGE = TRUE to convert invalid records to NULL automatically
- D. Set FORCE = TRUE to skip malformed records but keep valid rows

Answer: A

Correct answer: A. ON_ERROR controls how COPY behaves when parsing or conversion errors occur and should match the pipeline SLA. Why the strongest distractor fails: D. Set FORCE = TRUE to skip malformed records but keep valid rows — FORCE affects reloading files already recorded in load metadata; it does not control row-level parsing errors.

Question 24 : Domain 1.0: Data Movement

[Easy] After a failed test run, the media-streaming company revisits an ingestion pipeline with malformed rows. The engineer must reduce latency without hiding failures. Select the option containing the best primary action and the most relevant validation step.

- A. Set ON_ERROR according to the batch SLA, such as CONTINUE for quarantining bad rows while loading valid rows; then verify load history and rejected-row evidence after the run**
- B. Set FORCE = TRUE to skip malformed records but keep valid rows; then test recovery behavior within the documented retention boundaries
- C. Set PURGE = TRUE to convert invalid records to NULL automatically; then verify the persistence and recovery tradeoff matches the data classification
- D. Set MATCH_BY_COLUMN_NAME = CASE_INSENSITIVE to ignore all data quality errors; then document the recovery objective and confirm the object coverage

Answer: A

Correct answer: A. ON_ERROR controls how COPY behaves when parsing or conversion errors occur and should match the pipeline SLA. Why the strongest distractor fails: B. Set FORCE = TRUE to skip malformed records but keep valid rows; then test recovery behavior within the documented retention boundaries — FORCE affects reloading files already recorded in load metadata; it does not control row-level parsing errors.

Question 25 : Domain 1.0: Data Movement

[Easy] A senior data engineer at the payments processor is redesigning a cross-region data-distribution requirement. The design should support a production-grade design. The first proposed remediation targets the wrong Snowflake feature. Which change should the engineer recommend?

- A. Set FORCE = TRUE to skip malformed records but keep valid rows
- B. Set PURGE = TRUE to convert invalid records to NULL automatically
- C. Set MATCH_BY_COLUMN_NAME = CASE_INSENSITIVE to ignore all data quality errors
- D. Set ON_ERROR according to the batch SLA, such as CONTINUE for quarantining bad rows while loading valid rows**

Answer: D

Correct answer: D. ON_ERROR controls how COPY behaves when parsing or conversion errors occur and should match the pipeline SLA. Why the strongest distractor fails: A. Set FORCE = TRUE to skip malformed records but keep valid rows — FORCE affects reloading files already recorded in load metadata; it does not control row-level parsing errors.

Question 26 : Domain 1.0: Data Movement

[Easy] For the payments processor, a new requirement affects data-lake files that should remain in external storage. The implementation must support a production-grade design. Select the option containing the best primary action and the most relevant validation step.

- A. Create an external table over the staged files and refresh or automate metadata updates as required; then review ingestion history and retry behavior**
- B. Create a standard view that names the bucket URL directly in the FROM clause; then run a controlled restore or failover exercise
- C. Use a row access policy to discover newly added files; then test the policy or grant with authorized and unauthorized roles
- D. Create a task that clones the stage every minute; then inspect task or refresh history and validate downstream results

Answer: A

Correct answer: A. External tables let Snowflake query data in external storage while maintaining file metadata. Why the strongest distractor fails: B. Create a standard view that names the bucket URL directly in the FROM clause; then run a controlled restore or failover exercise — A standard view cannot use a bucket URL directly as a relational source.

Question 27 : Domain 1.0: Data Movement

[Easy] The IoT monitoring service is reviewing data-lake files that should remain in external storage. The priority is to make the smallest safe architectural change. The first proposed remediation targets the wrong Snowflake feature. Which implementation is most appropriate?

- A. Create a standard view that names the bucket URL directly in the FROM clause
- B. Use a row access policy to discover newly added files
- C. Create a task that clones the stage every minute
- D. Create an external table over the staged files and refresh or automate metadata updates as required**

Answer: D

Correct answer: D. External tables let Snowflake query data in external storage while maintaining file metadata. Why the strongest distractor fails: A. Create a standard view that names the bucket URL directly in the FROM clause — A standard view cannot use a bucket URL directly as a relational source.

Question 28 : Domain 1.0: Data Movement

[Easy] For the energy analytics program, a new requirement affects data-lake files that should remain in external storage. The implementation must reduce latency without hiding failures. The engineer must choose the option whose tradeoff matches the requirement. Which correction addresses the requirement most directly?

- A. Use a row access policy to discover newly added files
- B. Create a task that clones the stage every minute
- C. Create an external table over the staged files and refresh or automate metadata updates as required**
- D. Create a standard view that names the bucket URL directly in the FROM clause

Answer: C

Correct answer: C. External tables let Snowflake query data in external storage while maintaining file metadata. Why the strongest distractor fails: D. Create a standard view that names the bucket URL directly in the FROM clause — A standard view cannot use a bucket URL directly as a relational source.

Question 29 : Domain 1.0: Data Movement

[Easy] The media-streaming company is reviewing data-lake files that should remain in external storage. The priority is to support a production-grade design. A pull request uses a plausible but incorrect alternative. Which design is the best fit?

- A. Create a task that clones the stage every minute
- B. Create an external table over the staged files and refresh or automate metadata updates as required**
- C. Create a standard view that names the bucket URL directly in the FROM clause
- D. Use a row access policy to discover newly added files

Answer: B

Correct answer: B. External tables let Snowflake query data in external storage while maintaining file metadata. Why the strongest distractor fails: C. Create a standard view that names the bucket URL directly in the FROM clause — A standard view cannot use a bucket URL directly as a relational source.

Question 30 : Domain 1.0: Data Movement

[Easy] For the customer-support analytics team, a new requirement affects an inbound feed producing thousands of tiny files. The implementation must retain useful diagnostic evidence. The team is ready to implement one targeted change. Which approach should be selected?

- A. Disable compression so the warehouse can read more bytes
- B. Create one warehouse per file and load serially
- C. Consolidate very small source files into reasonably sized files before bulk loading**
- D. Split every source file into one row per file to maximize parallel operations

Answer: C

Correct answer: C. Reasonably sized files reduce metadata and scheduling overhead while still allowing parallel loading. Why the strongest distractor fails: B. Create one warehouse per file and load serially — One-row files create excessive overhead and usually reduce ingestion efficiency.

Question 31 : Domain 1.0: Data Movement

[Easy] The insurance data platform is reviewing an inbound feed producing thousands of tiny files. The priority is to use native Snowflake capabilities. Which statement identifies the correct Snowflake capability?

- A. Create one warehouse per file and load serially
- B. Consolidate very small source files into reasonably sized files before bulk loading**
- C. Split every source file into one row per file to maximize parallel operations
- D. Disable compression so the warehouse can read more bytes

Answer: B

Correct answer: B. Reasonably sized files reduce metadata and scheduling overhead while still allowing parallel loading. Why the strongest distractor fails: A. Create one warehouse per file and load serially — One-row files create excessive overhead and usually reduce ingestion efficiency.

Question 32 : Domain 1.0: Data Movement

[Easy] For the public-sector reporting unit, a new requirement affects an inbound feed producing thousands of tiny files. The implementation must preserve least-privilege access. The team wants a supported Snowflake pattern rather than a custom workaround. Which correction addresses the requirement most directly?

- A. Consolidate very small source files into reasonably sized files before bulk loading**
- B. Split every source file into one row per file to maximize parallel operations
- C. Disable compression so the warehouse can read more bytes
- D. Create one warehouse per file and load serially

Answer: A

Correct answer: A. Reasonably sized files reduce metadata and scheduling overhead while still allowing parallel loading. Why the strongest distractor fails: D. Create one warehouse per file and load serially — One-row files create excessive

overhead and usually reduce ingestion efficiency.

Question 33 : Domain 1.0: Data Movement

[Easy] The healthcare analytics team is reviewing an inbound feed producing thousands of tiny files. The priority is to retain useful diagnostic evidence. Select the option containing the best primary action and the most relevant validation step.

- A. Split every source file into one row per file to maximize parallel operations; then re-run representative workloads and examine queueing, scan, and spill metrics
- B. Disable compression so the warehouse can read more bytes; then validate that the measured bottleneck improves without an unjustified cost increase
- C. Create one warehouse per file and load serially; then document the recovery objective and confirm the object coverage
- D. Consolidate very small source files into reasonably sized files before bulk loading; then verify load history and rejected-row evidence after the run**

Answer: D

Correct answer: D. Reasonably sized files reduce metadata and scheduling overhead while still allowing parallel loading. Why the strongest distractor fails: C. Create one warehouse per file and load serially; then document the recovery objective and confirm the object coverage — One-row files create excessive overhead and usually reduce ingestion efficiency.

Question 34 : Domain 1.0: Data Movement

[Easy] A senior data engineer at the telecommunications provider is redesigning an inbound feed producing thousands of tiny files. The design should keep recovery behavior explicit. The team needs a maintainable solution that can be explained during an audit. Which change should the engineer recommend?

- A. Create one warehouse per file and load serially
- B. Disable compression so the warehouse can read more bytes
- C. Split every source file into one row per file to maximize parallel operations
- D. Consolidate very small source files into reasonably sized files before bulk loading**

Answer: D

Correct answer: D. Reasonably sized files reduce metadata and scheduling overhead while still allowing parallel loading. Why the strongest distractor fails: A. Create one warehouse per file and load serially — One-row files create excessive overhead and usually reduce ingestion efficiency.

Question 35 : Domain 1.0: Data Movement

[Easy] The manufacturing group is preparing a production rollout for an open-table-format requirement backed by cloud object storage. The team wants to preserve predictable behavior during retries. The chosen answer must address the measured bottleneck, not merely add components. Which change should the engineer recommend?

- A. Use a temporary table because temporary objects expose Iceberg metadata automatically
- B. Use Fail-safe as an open table catalog
- C. Use a materialized view to replace the external catalog and object storage
- D. Use an Iceberg table approach when interoperable open-table metadata and external cloud storage are explicit requirements**

Answer: D

Correct answer: D. Iceberg tables are relevant when the design requires interoperable open table formats backed by cloud storage. Why the strongest distractor fails: A. Use a temporary table because temporary objects expose Iceberg metadata automatically — A temporary Snowflake table does not provide Iceberg metadata interoperability.

Unlock All 1050+ Questions

Get the complete Q&A package with detailed explanations, topic analytics, and exam-accurate practice.

From €29.00

Visit: <https://cert-pass.com/exams/snowflake-snowpro-advanced-data-engineer>

CERT-PASS

cert-pass.com

© 2026 Cert-Pass. This material is for personal use only. Do not distribute.