

SAP C_BW4H

SAP Certified - Data Engineer - SAP BW/4HANA

For More Information – Visit link below:

<https://www.examsempire.com/>

Product Version

- 1. Up to Date products, reliable and verified.**
- 2. Questions and Answers in PDF Format.**



<https://examsempire.com/>

Visit us at: <https://www.examsempire.com/c-bw4h>

1. Unified Scenario Simulation

2.

Topic: 1

Unified Scenario Simulation

Aurora Grain Foods BW/4HANA Harvest Reporting Foundation

Business Context

Aurora Grain Foods is a mid-sized agri-food processor based in Winnipeg, Canada, milling and packaging oats, canola, and pulse products for retail and food-service customers across central Canada. Its planning and procurement teams want a single warehouse for harvest intake, milling yield, and distribution volumes, replacing a patchwork of spreadsheets that each region maintains separately. The analytics team has started building this warehouse on SAP BW/4HANA, the company's on-premise data warehouse powered by SAP HANA.

System Landscape

The data engineers are working through an early build. Intake volumes, supplier records, and plant master data come from the operational source through delta-capable extraction, while a small set of regional pricing references arrives as periodic flat files. The team has begun defining the characteristics and key figures that describe grain lots, plants, and suppliers, and is deciding which descriptions belong in reusable master data and which can stay as simple fields on a single report.

Description Modeling

A data engineer modeled grain-lot intake using characteristics and key figures, expecting plant and supplier descriptions to appear automatically on every report that uses them. During an early review, a colleague noticed that one analyst's quick report showed only technical codes for plants rather than readable names. The two reports were built differently: one reused shared master-data-bearing objects, while the other used plain fields defined only for that report. The same plants appeared correctly in the first report.

Acquisition Context

Intake data is loaded through transformations and a data transfer process from the delta-capable source. The team noticed that a first daily run brought in the full history, and a colleague asked whether later runs should repeat the whole load or carry only the new and changed records. Separately, the regional pricing file was loaded once during a rehearsal, and a reviewer flagged that the loaded values matched the file but the run was triggered manually rather than as part of any scheduled sequence.

Reporting and Operations Context

The reporting analyst wants planners to see milling yield by plant and month, with a measure that shows yield only for lots that passed quality screening. A first attempt summed all lots regardless of screening status. The operations lead has also asked how the daily intake load and the weekly pricing refresh will run without someone starting each step by hand, and how the team will confirm a load actually finished before planners open the morning reports.

Close Preparation

With the foundation taking shape, the team must settle which descriptions belong in reusable master data, confirm how the delta-capable source should behave after its first load, build the yield measure so it respects screening status, and decide how the daily and weekly loads will run and be confirmed. These early choices determine whether Aurora's planners open consistent, readable, trustworthy harvest reports each morning, or inherit gaps that surface only once the warehouse is in daily use.

Your assignment

You are acting as the data engineer supporting Aurora Grain Foods during the early build of its harvest-reporting warehouse on SAP BW/4HANA. Operational intake arrives from a delta-capable source, regional pricing arrives as periodic flat files, and planners need readable, screening-aware reports each morning. Your task is to advise on how descriptions should be modeled, how the delta-capable source

should behave after its first load, how the screening-aware yield measure should be built, and how routine loads should be scheduled and confirmed. You will reason about modeling and configuration decisions and recommend the soundest direction; you will not build objects directly

1. CHALLENGE 1 — Master Data Versus Field-Based Modeling Choice

- Advise on how plant and supplier descriptions should be modeled so their readable text is available to every consuming report rather than only one.
- Distinguish a reusable master-data-bearing object from a plain field defined for a single report.
- Reason about why one report showed readable plant names while another showed only technical codes for the same plants.

Checkpoints

- Confirm a report that reuses the shared object renders readable descriptions rather than technical codes.
- Decide which descriptive attributes warrant reusable master data versus a simple field.

2. CHALLENGE 2 — Delta Behavior After the Initial Source Load

- Advise on how the delta-capable source should behave after the first full load, so later runs carry only new and changed records.
- Reason about why repeating a full load on every run is undesirable for a delta-capable source.
- Recommend how the team confirms that subsequent runs are moving the delta rather than the full set.

Checkpoints

- Confirm that runs after the initial load transfer only new and changed records.
- Decide how to verify the delta is being applied rather than a repeated full load.

Question: 1

CHALLENGE 1 — Master Data Versus Field-Based Modeling Choice

Plant descriptions should appear readably on every report that uses the plant, not just one. Which modeling choice best achieves this across reports?

- A. Define the plant description as a plain field on each individual report so every analyst can type the names they prefer
- B. Store the plant description in a separate spreadsheet and paste it into each report after the data is loaded each day
- C. Leave the plant as a technical code only and let planners memorize which code maps to which physical plant location
- D. Model the plant as a master-data-bearing characteristic so its description is reusable across consuming reports

Answer: D

Explanation:

A master-data-bearing characteristic carries its descriptive text centrally, so every consuming report can display the readable plant name from the shared object. This delivers consistent, reusable descriptions instead of report-by-report duplication.

Question: 2

CHALLENGE 1 — Master Data Versus Field-Based Modeling Choice

One analyst's quick report showed only technical plant codes, while another report showed readable plant names for the same plants.

What most likely explains this difference?

- A. The readable report was loaded more recently, so its descriptions had simply not yet expired in the other report
- B. The code-only report used a plain field, while the readable report reused a master-data-bearing object
- C. The two reports queried different source systems, so only one source contained any plant descriptions at all
- D. The code-only report was opened in a different language, so its plant descriptions fell back to numeric codes

Answer: B

Explanation:

The readable report reused a master-data-bearing object that carries the plant description, while the code-only report used a plain field that has no associated text. The difference is the modeling choice, not the data freshness or source.

Question: 3

CHALLENGE 1 — Master Data Versus Field-Based Modeling Choice

The team wants to confirm the shared modeling choice actually delivers readable descriptions before standardizing on it.

Which validation best demonstrates this?

- A. Confirm the migration of the object completed without errors in the technical log for the latest build
- B. Confirm the daily intake load finished on schedule for the plants included in the report period
- C. Confirm a report reusing the shared object displays readable plant names rather than technical codes
- D. Confirm the pricing flat file loaded its values exactly as they appear in the source file rows

Answer: C

Explanation:

The modeling choice is validated by observing a consuming report render readable plant names from the shared object, which directly evidences the description is available across reports. This is the observable result the challenge targets.

Question: 4

CHALLENGE 2 — Delta Behavior After the Initial Source Load

After the first run brought in the full history, the team must decide how subsequent runs from the delta-capable source should behave.

Which behavior is correct for the runs that follow?

- A. Repeat the full history every run so the warehouse is always rebuilt completely from scratch each day
- B. Stop loading after the first run because the full history already covers everything the planners will ever need
- C. Switch the source off and re-enter new intake by hand each day so the team keeps tight control of every record
- D. Carry only the new and changed records as a delta, since the delta-capable source supports incremental loading

Answer: D

Explanation:

A delta-capable source supports moving only new and changed records after an initial full load, which is the efficient and correct steady-state behavior. Subsequent runs should apply the delta rather than reload everything.

Question: 5

CHALLENGE 2 — Delta Behavior After the Initial Source Load

A reviewer wants to confirm that runs after the initial load are behaving as intended.

Which check most directly confirms correct delta behavior?

- A. Confirm later runs transfer only new and changed records rather than the full history
- B. Confirm the readable plant descriptions appear on the planning report for the loaded period
- C. Confirm the yield measure restricts its result to lots that passed quality screening
- D. Confirm the weekly pricing flat file matches the values in the source file exactly

Answer: A

Explanation:

Correct delta behavior is confirmed by observing that runs after the initial load move only new and changed records rather than reloading the full history. This directly evidences the incremental load is working.

Topic: 2

Micro Skill Drill Exam

Question: 6

Fjordmat, an Oslo-based grocery retailer, loads daily point-of-sale figures from a source system into a DataStore Object (advanced) using a transformation and a Data Transfer Process. A field-level rule in the transformation derives the net sold quantity for each article. After last night's run, the load monitor shows the request completed successfully and new records are present in the target for today's stores. However, analysts report that the net quantity column is zero for every new record, while article, store, and date values populated correctly. The administrator confirms the request activated, no errors or red entries appear in the log, and prior historical records still hold their quantities. Nothing in the run flags a

failure; only the single quantity field is empty. Working in the modeling environment, the administrator must explain why a clean, successful load delivers zero into one field and apply the correct fix before the daily report is published.

What is the root cause of the zero quantity values, and how should it be corrected?

- A. The source DataSource excludes the quantity field from its extraction structure, so no value is delivered to the transformation; extend the DataSource so the field reaches the target.
- B. The transformation rule for the quantity field is not assigned in the active rule group, so the target receives its initial value; map the source field to the quantity rule so loads populate it.
- C. The SAP HANA delta merge has not yet run on the DataStore Object, so the loaded quantities remain unmerged and read as zero; trigger the delta merge so the values become visible.
- D. The Data Transfer Process executed in delta mode before any initialization, so only empty deltas reached the target; switch the process to full extraction once to backfill the missing quantity values.

Answer: B

Explanation:

When a field's transformation rule is absent from the active rule group, the target field is written with its initial value — zero for a quantity — even though the request completes cleanly and all mapped fields populate. This matches the evidence exactly: a green load, correct article, store, and date, and one empty field. Restoring the mapping for the quantity rule lets each load derive and write the value.

Question: 7

Companhia Energia Sul, a regional electricity utility in São Paulo, publishes regional consumption analytics from SAP BW/4HANA. A new analyst on the Southeast desk is supposed to see consumption only for the Southeast region, in line with the utility's data-access policy. The modeler added the analyst to a role that grants execution of the regional consumption query, and the analyst can open and run it without error. During review, a supervisor notices the analyst's results list every region's consumption, not just the Southeast. The query itself is unchanged and works correctly for other users whose visibility is already scoped. The modeler confirms the analyst has no special administrative profile and that the data model and query definition are the same ones used across the desk. The team must explain why an apparently correct access setup still exposes all regions and apply the change that enforces the intended per-region restriction.

Why does the analyst see all regions despite the restriction intent, and what enforces the correct scope?

- A. The query's region variable defaults to all values, so set a mandatory region variable on the query so the analyst is prompted for a single region at runtime.
- B. The DataStore Object does not carry region as a key field, so add region to the object key so that the stored data can be filtered down to the analyst's permitted region.
- C. The analyst holds all-data access through an inherited role, so remove the query-execution rights from that role so the regions the analyst can see reduce to just the one assigned.
- D. Standard authorizations grant execution access, not data-level scoping; add an analysis authorization on the region characteristic so the analyst sees only the Southeast region.

Answer: D

Explanation:

Standard authorizations control which objects and queries a user may execute, but they do not restrict which data values are returned. Region-level visibility is data-level security, which BW/4HANA enforces through analysis authorizations on the relevant characteristic. Creating an analysis authorization that limits the region characteristic and assigning it returns only the permitted region while preserving query execution.

Question: 8

Murray Plains Agri, an agribusiness in Adelaide, needs a single report combining current-season and historical harvest volumes that currently live in two separate advanced DataStore Objects. A modeler, aiming to deliver quickly, loaded both objects' data into one new standard DataStore Object and pointed the report at it. The report runs, but season totals are visibly inflated: volumes for periods present in both sources appear roughly doubled, while periods unique to one source look correct. The modeler confirms each source object individually reports accurate volumes and that no transformation changes the figures during the copy. Business users need accurate combined totals without maintaining a third physical copy of the data. The team must explain why the combined report double-counts and choose the modeling approach that combines the two sources correctly for reporting.

Why does the combined report double-count volumes, and what is the correct modeling choice?

- A. The Data Transfer Process aggregates records during load, so disable its aggregation setting so the overlapping season volumes are summed only once at report time instead of twice.
- B. The calculated key figure uses summation aggregation, so change it to maximum so the duplicated volumes stop adding together in the report.
- C. An Open ODS View exposes raw rows without keys, so convert it to an InfoObject-based model so the duplicate records collapse on the key.
- D. Loading both sources into one standard DataStore Object merges overlapping records and inflates totals; use a CompositeProvider to union the two providers virtually instead.

Answer: D

Explanation:

Physically loading both sources into one standard DataStore Object combines records for periods present in both, so overlapping volumes accumulate and totals inflate. Modeling a CompositeProvider that unions the two providers combines them virtually for reporting without copying data, so each record is counted once and no duplication is introduced.

Question: 9

Boreal Potash, a mining group in Saskatchewan, runs a nightly process chain that loads several advanced DataStore Objects and then refreshes an executive CompositeProvider report. This morning the report shows the prior day's figures even though the overnight job finished and the chain is entirely green. The administrator confirms the source delivered new records, the load requests are present, and no step turned red. Other reports that read different objects are current. Inspecting the chain, the administrator sees that the step which makes the newly loaded data available to reporting for one key DataStore

Object and the step that builds the executive report are both present, but the report appears to consume the object before the day's new data has been made active. The team must explain why a green chain yields a stale report and correct the chain so the executive view reflects the current load. Why is the report stale despite a successful, green process chain, and how should the chain be corrected?

- A. The delta merge step is disabled in the chain, so SAP HANA keeps the records unmerged and the report cannot read them; enable the delta merge step so the data becomes visible.
- B. The activation step that moves loaded data to the active table runs after the report consumes the object, so reorder the chain to activate before the report reads it.
- C. The Data Transfer Process loads in full mode, so each run overwrites the target; switch it to delta so the day's new records reach the report.
- D. The CompositeProvider caches its results between runs, so add a step to clear its cache after the load so the report displays the latest figures.

Answer: B

Explanation:

An advanced DataStore Object makes loaded data available to reporting only after its activation step moves records to the active table. If the chain consumes the object for the report before that activation runs, the report reads the prior active state and appears stale despite a green chain. Reordering the chain so activation precedes the reporting step makes the current data available.

Question: 10

Deccan Freightways, a logistics provider in Pune, modernized its data acquisition by re-pointing shipment extraction from an older mechanism to an ODP-based delta feed from its source system. Since cutover, the daily delta loads complete successfully but bring noticeably fewer shipment records than the business expects, and downstream reports under-report volumes. When the team runs a full extraction for the same period, the record counts are correct and match the source. The daily deltas are clearly running — requests appear each day with some records — but the change set is consistently short. The transformation and Data Transfer Process are the same ones validated during testing, and no errors are raised. The team must explain why daily deltas under-report while a full load is accurate, and identify the corrective action so ongoing deltas capture the complete set of changes.

Why do the daily deltas under-report while a full load is accurate, and what is the correct corrective action?

- A. The delta was not initialized correctly against the ODP source after re-pointing, so it starts from an incomplete baseline; run a proper delta initialization to capture all changes.
- B. The Data Transfer Process is set to full mode and overwrites rather than appends, so switch it to delta mode so the daily shipment records accumulate over time instead of being replaced on each run.
- C. An end routine in the transformation filters out shipment records, so remove the filter so the daily deltas pass the missing records through to the target.
- D. The SAP HANA delta merge lags on the target, so schedule it earlier so that the daily shipment records appear sooner in the downstream operational reports.

Answer: A

Explanation:

When extraction is re-pointed to the ODP source without a correct delta initialization, the delta captures only changes after an incomplete starting point, so daily loads under-report while a full extraction still reads everything. Performing a proper delta initialization re-establishes the baseline so subsequent deltas capture the complete change set.

Thank You for Trying Our Product

Special 16 USD Discount Coupon: NSZUBG3X

Email: support@examsempire.com

**Check our Customer Testimonials and ratings
available on every product page.**

Visit our website.

<https://examsempire.com/>