SAP
C_BW4H_2601
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.**



Latest Version: 4.0

1. Micro Skill Drill Exam
2. Unified Scenario Exam

Topic: 1
Micro Skill Drill Exam

Question: 1

Nagpur Metro Transit, an urban rail operator in Nagpur, runs a process chain that loads ridership data on two parallel branches — one for gate-entry counts and one for ticket-revenue records — before a final step builds a daily operations report that combines both. This morning the report shows complete gate-entry counts but missing ticket revenue, even though the chain is green and both branches eventually finished. The administrator confirms both source feeds delivered data, both load requests are present and activated, and the report step itself works when run manually afterward. Inspecting the chain, the administrator sees the two branches feed into a collector step that triggers the report, and that the collector releases the report as soon as the first branch completes rather than waiting for both. Because the gate-entry branch usually finishes first, the report runs before ticket revenue has loaded. The team must explain why the report omits ticket revenue despite a green chain and correct the chain so it combines both branches every day.
Why does the report omit ticket revenue despite a green process chain, and how should the chain be corrected?
Response:

A. The collector step releases the report when any one branch finishes, so it runs before ticket revenue loads; change the collector so the report is triggered only after both branches complete.
B. The ticket-revenue branch loads in delta mode while the report expects full data, so revenue is missing; switch that branch to full mode so the report sees complete revenue figures each day.
C. The activation step for the ticket-revenue records runs after the report step, so revenue is not yet active; move the activation earlier so the report reads activated revenue when it builds.
D. The report reads a CompositeProvider that caches the prior run, so revenue appears missing; add a step to clear the cache before the report builds so current revenue is included.

Answer: A

Explanation:
Feedback:
A collector that releases its successor as soon as any one predecessor finishes will trigger the report after whichever branch completes first. With the gate-entry branch typically finishing ahead of ticket revenue, the report builds before revenue has loaded, so the chain is green yet the report is incomplete. Configuring the collector so the report is triggered only after both branches complete ensures both data sets are present every day.

Question: 2

Visit us at: https://www.examsempire.com/c-bw4h-2601

Jadran Maritime Logistics, a port-logistics operator in Zagreb, tracks each container's status as it moves through the terminal. The source sends a fresh status record for a container several times a day as it progresses, always keyed by the container number. A modeler loaded these updates into a DataStore Object (advanced) modelled so that records sharing the same key overwrite one another, then pointed an operations report at it. The report's current-status view is accurate, but a new requirement to analyse dwell time between status changes returns almost nothing, because only the latest status per container is retained. The modeler confirms the source genuinely sent every intermediate status and that no transformation drops records. The team must explain why intermediate statuses are missing and choose the modelling approach that preserves the full sequence of changes for analysis.
Why are the intermediate statuses missing, and what is the correct modelling choice?
Response:

A. The Data Transfer Process is filtering out intermediate statuses, so widen its selection so every status record for each container is loaded into the object instead of only the latest.
B. The source omits intermediate statuses on delta runs, so initialise the delta again so the object receives the full set of status changes for each container.
C. The object overwrites records that share the container key, so each new status replaces the prior one; model it to capture changes so every status is retained as its own record.
D. The SAP HANA delta merge collapses repeated container keys during consolidation, so disable the merge so the intermediate statuses are no longer combined into one row.

Answer: C

Explanation:
Feedback:
When a DataStore Object is modelled so that records sharing a key overwrite, each new status for a container replaces the previous one, leaving only the latest — accurate for current status but unable to support change analysis. Modelling the object to capture changes keeps every delivered status as its own record, so the full sequence is retained and dwell time between changes can be analysed.

Question: 3

Têxtil Serra Gaúcha, a textile manufacturer in Caxias do Sul, Brazil, loads daily sales into an advanced DataStore Object. A transformation reads a customer master attribute — the customer's assigned sales region — and writes it onto each sales record so reports can analyze sales by region. A nightly process chain loads the sales records, and a separate weekly chain refreshes the customer master. This week several new customers were onboarded and began ordering immediately. Their sales loaded green into the target, but their region is blank, so their revenue lands in an unassigned bucket; established customers carry their region normally. The administrator confirms the new customers exist in the source, their master data was created, and the transformation rule that derives the region is correct and unchanged. The team must explain why correctly modeled records still load without a region for brand-new customers, and apply the corrective action so newly onboarded customers report under their region.
Why do new customers' sales load with a blank region despite a correct derivation rule, and what is the correct corrective action?
Response:

A. The derivation rule lost its mapping in the active rule group, so it writes the initial value for every record on the load; reassign the region rule in the rule group so the region once again populates on each and every record loaded.
B. The transformation looks up the region from customer master that the weekly chain has not yet refreshed, so new customers have no master entry to read at sales-load time; load the customer master before the sales load so the region resolves.
C. The new customers exceed the analysis authorization for the region characteristic, so the region is suppressed on their rows; extend the authorization to the new customers so their region becomes visible in the report.
D. The Data Transfer Process ran in delta mode before initialization for the new customers, so their region arrived empty from the source system; run a full load once for the new customers so their missing region values backfill correctly into the target.

Answer: B

Explanation:
Feedback:
The transformation derives the region by reading the customer master, so the value depends on the master being present when the sales record is processed. Because the customer master is refreshed on a separate weekly cycle, customers onboarded after the last refresh have no master entry yet, and the lookup returns nothing while the sales load still completes green. Loading or refreshing the customer master before the sales load gives the rule a record to read, so the region resolves for the new customers.

Question: 4

Cebu Island Resorts, a hospitality group in the central Philippines, runs an occupancy query in SAP BW/4HANA with a restricted key figure showing revenue for one property cluster, defined by restricting the property characteristic to a specific node of the property hierarchy. The query was accurate for the prior year. After the organization restructured its property hierarchy and renamed and regrouped its clusters, finance reports that the cluster revenue column is now empty for every row, while the overall revenue column and all other characteristics still display correctly. The administrator confirms the current data is loaded and visible, no authorization message appears, and the hierarchy itself is active and used elsewhere without issue. Reviewing the query, the administrator sees the restricted key figure still points at the node identifier that existed before the restructure. Nothing else in the query changed. The team must explain why the cluster column emptied exactly after the hierarchy was reorganized, and apply the design correction so the column resolves to the intended cluster going forward.
Why did the cluster revenue column go empty after the hierarchy was reorganized, and what is the correct design fix?
Response:

A. A calculated key figure divides cluster revenue by an occupancy base that is now zero, so the column returns empty; add a guard for the zero denominator so the values render again.
B. The restricted key figure points at a hierarchy node that the restructure removed or regrouped, so its restriction now matches nothing; repoint it to the current node so the cluster resolves.

C. An analysis authorization for the reorganized cluster is missing, so its rows are suppressed for the buyer; grant the new node so the cluster revenue becomes visible again in the column for the report.
D. The hierarchy is no longer assigned to the property characteristic in the query, so nodes do not resolve; reassign the hierarchy so the characteristic resolves its nodes and the cluster column populates.

Answer: B

Explanation:
Feedback:
A restricted key figure that pins a specific hierarchy node depends on that node continuing to exist as defined. When the hierarchy is reorganized and the node is removed or regrouped, the restriction matches no members and the column empties, even though overall revenue and other characteristics remain correct because they do not depend on that node. Repointing the restriction to the current node — or to a variable that resolves it — restores the cluster figure and keeps it aligned after future changes.

Question: 5

Basel Diagnostics, a Swiss medical-laboratory group, loads results into an advanced DataStore Object and then feeds a downstream reporting object by extracting changes from the first object using a delt a. For months this worked, but after a recent housekeeping action the downstream delta loads now bring no new records, while the first DataStore Object itself clearly holds the latest results and reports correctly. The administrator confirms the first object is current, the downstream transformation and Data Transfer Process are unchanged, and the delta requests run green with zero records. A full load from the first object to the downstream object does deliver the data correctly. The housekeeping action recently cleared change history on the first object to reclaim space. The team must explain why the downstream delta now captures nothing while a full load works, and identify the corrective action.
Why does the downstream delta capture nothing while a full load works, and what is the correct corrective action?
Response:
A. The downstream object was modeled as a staging type with no active table, so the incoming delta records are discarded rather than stored; remodel it as a reportable type so the loaded data is retained.
B. The first object's analysis authorization no longer permits the downstream load, so grant the authorization so the delta records pass through.
C. The first object's calculated figures changed aggregation, so the delta reads zero amounts; reset the aggregation so the downstream delta carries values.
D. The change log the downstream delta reads was cleared by housekeeping, so the delta has no recorded changes to extract; reinitialize the delta to set a new baseline.

Answer: D

Explanation:
Feedback:
A delta from a DataStore Object reads the recorded changes in that object's change log; if housekeeping clears that history, there are no recorded changes for the next delta to extract, so it runs green with zero records even though the source data is current and a full load still works. Reinitializing the delta establishes a fresh baseline so subsequent change-based loads capture new records again.

Question: 6

Yarra Health Devices, a medical-device maker in Melbourne, feeds a reporting DataStore Object (advanced) from an upstream advanced DataStore Object using a delta Data Transfer Process. The upstream object receives daily production records correctly, and a report directly on it is always current. The downstream reporting object, however, stays empty after each daily run: the Data Transfer Process finishes green and reports zero records transferred, day after day. The administrator confirms the upstream object holds the day's new records, the transformation between the two is valid, and no errors appear. A one-time full transfer between the same two objects does move all records correctly. Only the recurring delta brings nothing. The team must explain why the delta consistently transfers zero records from an upstream object that clearly holds new data and identify the corrective action so the daily delta carries the changes forward.
Why does the delta Data Transfer Process transfer zero records while a full transfer works, and what is the corrective action?
Response:

A. The upstream object's change log is not maintained, so the delta has no record of changes to extract; enable change-log maintenance so the delta can read each day's new records.
B. The delta was never initialized between the two objects, so it has no baseline; run the initialization so the recurring delta picks up subsequent changes.
C. The downstream object is a staging type with no active data, so transferred records are discarded; remodel it with a type that retains active data so the delta's records persist.
D. The transformation drops all rows during the delta path, so add a rule group for the delta so the records are mapped and transferred.

Answer: A

Explanation:
Feedback:
A delta transfer between advanced DataStore Objects sources its changes from the upstream object's change log; if that log is not maintained, no changes are recorded for the delta to read, so it consistently transfers zero even though the active data holds the day's records. A full transfer reads the active data directly, which is why it works. Enabling change-log maintenance gives the delta the change records it needs each day.

Question: 7

Abidjan Cocoa Exchange, a commodities trading house in Côte d'Ivoire, exposes intraday position data in SAP BW/4HANA through an Open ODS View that reads directly from a source system, so traders always see live figures. A new end-of-day reconciliation report was built on the same Open ODS View. Auditors find the reconciliation figures differ each time they are rerun for the same closed day, and they do not match the snapshot the business agreed at cutoff. The administrator confirms the Open ODS View is correctly defined, the source is reachable, and intraday users are satisfied with the live behavior. The reconciliation, however, needs a stable, repeatable end-of-day figure that does not change after cutoff.

The team must explain why the reconciliation figures keep shifting and choose the modeling approach that gives a stable end-of-day basis without disturbing the live intraday view.
Why do the reconciliation figures keep shifting, and what modeling approach gives a stable basis?
Response:

A. The Open ODS View caches stale rows between runs, so add a step to clear its cache after cutoff so each reconciliation run returns the agreed figure.
B. The reconciliation reads the Open ODS View's live virtual data, which keeps changing; persist the end-of-day data into a DataStore Object for a stable basis.
C. An analysis authorization restricts which rows the reconciliation sees, so grant the missing values so every run returns the same complete figure after cutoff.
D. The Data Transfer Process loads in delta mode, so reconciliation sees only changes; switch it to full extraction so the closed day's figures stay constant on rerun.

Answer: B

Explanation:
Feedback:
An Open ODS View provides virtual, live access, so a report on it reflects the source at the moment it runs and changes whenever the source changes after cutoff. Reconciliation needs a figure captured and frozen at end of day, which means persisting that data into a DataStore Object. Loading the closed day into a persisted object gives a stable, repeatable basis while the intraday view keeps reading live.

Question: 8

Wuhan Optics Manufacturing publishes a SAP BW/4HANA query that lists suppliers with two columns: a returns column built as a restricted key figure showing only returned-goods value, and a separate total purchases column. The report has zero suppression switched on to keep the list compact. Procurement reports that many active suppliers are missing from the list entirely, even suppliers with substantial purchases, while suppliers that do appear show correct figures in both columns. The administrator confirms the data is loaded, the missing suppliers have purchase activity in the period, no authorization message is raised, and the report shows every supplier when zero suppression is turned off. Looking closer, the administrator sees that the suppliers who vanish are exactly those with no returns in the period, so their returns column is zero while their purchases column is not. The team must explain why suppliers with real purchase activity drop off the list, and apply the correction so every supplier with purchases appears while the list stays readable.
Why do suppliers with real purchase activity drop off the list, and what is the correct fix?
Response:

A. The returns restricted key figure points at a value no supplier carries, so update its restriction so the returns column populates and the dropped suppliers reappear on the list.
B. The purchases column is loaded in delta mode and misses some suppliers, so reload it in full mode so the missing suppliers' rows return to the supplier list.
C. Zero suppression drops rows that evaluate to zero across the suppressed scope, so adjust it to keep a supplier whose purchases are non-zero even when returns is zero.
D. An analysis authorization restricts the missing suppliers, so grant the values so they appear in the list with their purchase and returns figures shown.

Answer: D

Explanation:
Feedback:
Zero suppression removes rows whose figures evaluate to zero across the scope it is applied to, so when a supplier's returns column is zero it can be dropped even though its purchases column carries a value, which matches active suppliers with no returns disappearing while the data is intact. Adjusting zero suppression so a row is kept when any displayed figure such as purchases is non-zero, rather than dropping it on the zero returns column, keeps every purchasing supplier on a still-compact list.

Question: 9

Vestjysk Vind, a wind-energy firm in Aarhus, reports turbine generation in SAP BW/4HANA where the turbine characteristic carries a site attribute identifying the wind farm each turbine belongs to. Analysts can see the site name beside each turbine in the report, but they cannot pivot generation by site: the site does not appear as a characteristic they can add to the rows or drill into, so they cannot view totals per wind farm. The administrator confirms the data is loaded, the attribute value is populated for every turbine, and the query is otherwise correct. Reviewing the InfoObject, the administrator sees the site attribute was modeled so it can be shown next to the turbine but not used for navigation. The team must explain why the site can be displayed but not used to drill down, and apply the modeling correction so analysts can pivot generation by wind farm.
Why can analysts see the site attribute but not drill down by it, and what is the correct modeling fix?
Response:

A. The site is set as a display attribute, not a navigation attribute, so it shows but cannot be pivoted; switch it to navigation so analysts can drill by site.
B. An analysis authorization omits the site values, so the characteristic is hidden from drilldown; grant the site values so analysts can add it to the rows and view totals.
C. The site attribute is missing from the DataStore Object key, so it cannot group records; add it to the key so the report can total generation per wind farm.
D. A restricted key figure pins a single site, so the others are suppressed from drilldown; widen the restriction so all wind farms appear when analysts pivot generation by site.

Answer: A

Explanation:
Feedback:
An attribute modeled for display can be shown beside its characteristic but cannot be placed in the rows or used to drill, which is exactly what the analysts experience. Enabling the attribute for navigation makes it available as a characteristic in the query so users can add it, pivot, and view generation totals per wind farm, while the value already populated continues to display.

Question: 10

Swan River Brewing, a craft brewery in Perth, publishes a sales-by-product-group report from SAP BW/4HANA. Management expects the product characteristic to appear in the report grouped under its product-group structure, so totals roll up by group with individual products beneath each group. A modeler maintains an active product-group structure for the characteristic, and other reports show it grouped correctly. In this query, however, the product characteristic lists every product in a single flat sequence with no grouping, even though every product value displays and the figures are accurate. The administrator confirms the structure exists, is active, and contains the products in question, and that no authorization message appears. Nothing in the query is broken; the numbers are right, but the requested grouped presentation is missing. The team must explain why the products appear flat rather than grouped and apply the query-design correction so the report presents the products under their group structure.
Why does the product characteristic display as a flat list instead of grouped, and how should the query be corrected?
Response:

A. The product master data was loaded after the transaction data, so the grouping links are missing; reload the master data so the products attach to their groups in the report.
B. The query is not set to present the product characteristic using its structure, so set it to use the active group structure so products appear under their groups.
C. A condition is filtering the query to top products only, so the grouping cannot form; remove the condition so the full grouped structure can be displayed.
D. The group structure is assigned but not active in this client, so activate it so the report can read the grouping and present the products beneath their groups.

Answer: B

Explanation:
Feedback:
A characteristic renders grouped only when the query is set to present it using its active structure; if that presentation is not selected, the characteristic shows a flat list of values even though the structure exists and the figures are correct. This matches the evidence exactly. Setting the characteristic to display using the group structure makes the products appear beneath their groups with rolled-up totals.
Topic: 2
Unified Scenario Exam

Question: 11

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?
Response:

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:
Feedback:
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: 12

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?
Response:

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:
Feedback:
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: 13

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?
Response:

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:
Feedback:
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: 14

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?
Response:

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:
Feedback:
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: 15

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?
Response:

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:
Feedback:
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.

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/