

# SAP

## C\_HAMOD\_2601


### SAP Certified - Data Engineer - SAP HANA

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-hamod-2601>

# Latest Version: 4.0

1. Micro Skill Drill Exam
2. Unified Scenario Exam

**Topic: 1**  
**Micro Skill Drill Exam**

## Question: 1

A regional bakery franchise builds an SAP HANA Cloud calculation view to analyze production yield. The view combines baked-unit counts, planned-unit counts, and store-format attributes. Activation succeeds, but preview shows correct yield percentages at bakery-shift level and misleading percentages when users aggregate by store format.

The source counts are validated, and no remote source is involved. The model must support both shift-level investigation and format-level monitoring from the same reusable artifact. The team must avoid dashboard-specific formulas and preserve the current source tables. The correction must address how the calculated percentage behaves during aggregation.

Which action best stabilizes the production-yield percentage?

Response:

- A. Persist the bakery-shift yield percentage in the source table so the calculation view reads stored values.
- B. Add a mandatory prompt requiring users to select bakery-shift detail before yield percentages are displayed.
- C. Create a separate calculation view for store-format monitoring so aggregation does not reuse shift-level results.
- D. Review the percentage measure definition, aggregation behavior, and model grain so the value is calculated from aggregated baked and planned unit counts.

**Answer: D**

Explanation:

Feedback:

This targets the semantic modeling layer where calculated measures must align with analytical grain. When bakery-shift values are correct but store-format percentages are misleading, the view must derive the measure from aggregated baked and planned unit counts at the query level.

## Question: 2

A private clinic network prepares appointment billing data in SAP HANA Cloud before exposing it through calculation views. After a transformation change, regional analysts can preview appointment totals, but one authorized region sees fewer billable appointments than expected. A developer with broader privileges sees the full transformed result, and the source appointment extract contains the missing records.

The environment uses modeled runtime access controls. The team must preserve regional data protection rules and avoid granting broad access for validation. The correction must ensure that authorized users see all billable appointments within their permitted scope.

Which action best identifies and corrects the issue?

Response:

- A. Disable regional restrictions temporarily so analysts can confirm that all transformed appointment records exist.
- B. Add missing appointment rows through a transformation rule so the regional preview returns the expected totals.
- C. Create a separate calculation view for the affected region with fewer access restrictions for validation.
- D. Review the modeled access scope, regional value mapping, and runtime user assignment applied to the transformed appointment dataset.

**Answer: D**

Explanation:

Feedback:

This targets the security and execution path that determines which transformed rows a regional user can retrieve. When broader privileges expose the full result and a regional user sees fewer permitted rows, the likely correction is in access scope mapping or runtime assignment rather than data transformation.

### Question: 3

A performing arts venue builds an SAP HANA Cloud calculation view to analyze ticket utilization. The view combines issued ticket counts, scanned entry counts, and seating section attributes. Activation succeeds, but preview shows correct utilization at event-section level and misleading utilization percentages when users aggregate by performance series.

The source counts are validated, and no remote source is involved. The model must support both event-level investigation and series-level monitoring from the same reusable artifact. The team must avoid dashboard-specific formulas and preserve the current source tables. The correction must address how the modeled percentage behaves during aggregation.

Which action best stabilizes the utilization percentage?

Response:

- A. Persist the event-section utilization percentage in the source table so the calculation view only reads stored values.
- B. Create a separate calculation view for performance-series monitoring so aggregation does not reuse event-section results.
- C. Review the utilization measure definition, aggregation behavior, and model grain so the percentage is calculated from aggregated scanned and issued ticket counts.
- D. Add a prompt requiring users to select event-section detail before utilization percentages are displayed.

**Answer: C**

Explanation:

Feedback:

This targets the semantic modeling layer where calculated measures must align with analytical grain. When event-section values are correct but series-level percentages are misleading, the view must calculate utilization from aggregated scanned and issued ticket counts at the query level.

## Question: 4

A food safety analytics team uses SAP HANA Cloud virtual tables to access supplier certification data from a remote quality system. A calculation view joins local batch inspection facts with the virtual certification table. Activation succeeds, but preview returns blank certification expiry dates for suppliers whose certifications were recently renewed.

The virtual table preview shows the renewed expiry dates for the same supplier IDs. The remote quality system must remain authoritative, and the team must preserve virtualized access. The correction must avoid copying certification records locally or changing validated inspection facts.

Which action best addresses the likely modeled retrieval dependency?

Response:

- A. Replace blank expiry dates with a default future date in the calculation view so validation can continue.
- B. Replicate supplier certification records locally so renewed expiry dates no longer depend on remote access.
- C. Validate the virtual table metadata, supplier key mapping, and calculation view binding used when the modeled join executes.
- D. Remove certification expiry dates from the calculation view and require users to query the remote source separately.

**Answer: C**

Explanation:

Feedback:

This addresses the dependency between remote metadata, key-field mapping, and calculation view binding. When the virtual table preview is current but the joined model returns blanks, the correction must verify how the model interprets and binds the current virtual structure during execution.

## Question: 5

A media streaming company prepares subscription usage data in SAP HANA Cloud before exposing it through a calculation view. After a transformation adjustment, premium usage events are correctly loaded, but some trial-account events appear in the paid revenue dataset. The source extract contains an account type indicator, and the consuming calculation view aggregates the transformed dataset without errors.

The environment uses development tooling and runtime validation of modeled artifacts. The team must keep the shared revenue model stable for finance dashboards and avoid placing account-type filters in

each consuming report. The correction must ensure that only eligible revenue events reach the calculation view.

Which action best corrects the issue at the proper dependency layer?

Response:

- A. Add a report-level account type filter so finance dashboards exclude trial-account events during consumption.
- B. Correct the transformation eligibility rule so trial-account events are excluded before the revenue dataset is consumed by the calculation view.
- C. Change the calculation view aggregation behavior so trial-account usage events contribute zero revenue when grouped.
- D. Remove the account type indicator from the source extract so the transformation no longer receives mixed event categories.

**Answer: B**

Explanation:

Feedback:

This acts at the transformation layer where event eligibility is determined before analytical modeling. Since the calculation view correctly aggregates what it receives, the correct correction is to ensure only eligible paid usage events are passed into the revenue dataset.

### Question: 6

A regional marine fuel supplier uses SAP HANA Cloud virtual tables to access vessel classification data from a remote operations source. A calculation view joins local fuel delivery facts with the virtual classification table. Activation succeeds, and the virtual table preview is fast, but calculation view preview times out when analysts filter by vessel class for a single reporting period.

The remote operations source must remain authoritative, and the team must preserve virtualized access. The filter condition is applied in the calculation view, but runtime evidence suggests that a large remote dataset is retrieved before the selective condition is applied. The correction must avoid local replication and keep the existing analytical model structure reusable.

Which action best addresses the second-order execution cause of the timeout?

Response:

- A. Increase the preview timeout so the calculation view has enough time to retrieve and join the remote vessel classification dataset.
- B. Replicate vessel classification data locally so the filter no longer depends on remote retrieval during calculation view execution.
- C. Add a dashboard warning that vessel-class filtering may be slow when the remote source contains many records.
- D. Validate filter placement, join-field compatibility, and virtual table binding so the selective condition can restrict remote retrieval before the modeled join expands execution.

**Answer: D**

Explanation:

Feedback:

This targets the execution dependency between filter application, remote retrieval, and modeled join binding. When a virtual table is fast alone but the joined model times out under a selective filter, the calculation view must be validated so the restriction is applied at the correct stage of execution.

## Question: 7

A book retailer models online order conversion in SAP HANA Cloud. A calculation view combines web session counts, completed order counts, and marketing channel attributes. Activation succeeds, but preview shows a conversion rate that changes incorrectly when users drill from campaign level to marketing channel level. The underlying session and order counts are validated.

The model must support campaign analysis and channel-level monitoring from the same reusable artifact. The team must avoid hardcoding formulas in consuming dashboards or creating separate views for each analytical level. The correction must preserve the intended conversion-rate semantics.

Which action best stabilizes the conversion-rate calculation?

Response:

- A. Persist the conversion rate in the source table so the calculation view no longer computes it during preview.
- B. Add a mandatory prompt that forces users to select either campaign level or channel level before viewing the rate.
- C. Create one calculation view for campaign analysis and another calculation view for channel monitoring.
- D. Review the calculated measure definition, aggregation behavior, and attribute relationship grain so conversion rate is derived from aggregated sessions and orders at query level.

**Answer: D**

Explanation:

Feedback:

This targets the modeling layer where base measures and attributes determine calculation semantics.

When validated counts produce unstable rates across drill levels, the conversion rate must be calculated from aggregated sessions and orders at the level being queried.

## Question: 8

A regional ferry operator provisions passenger boarding records into SAP HANA Cloud for daily capacity analytics. Initial boarding records arrive in the target table, but later source-side corrections to passenger category remain unchanged in the replicated dataset. The replication task reports completion without errors, and the calculation view consuming the target table activates successfully.

Operations analysts can count total boardings, but category-level reporting remains inaccurate for corrected records. The team must keep the existing replicated target because multiple operational models consume it. The correction must avoid dashboard-level recoding and identify why later category changes are not applied after the initial record insert.

Which action best addresses the likely provisioning-layer issue?

Response:

- A. Add a calculation view expression that derives passenger category from fare amount during analytical preview.
- B. Manually update passenger categories in the target table before each daily reporting cycle.
- C. Validate that the replication configuration captures passenger category update events and applies them to existing target rows.
- D. Create a separate replicated table for passenger category corrections and require consuming models to join it.

**Answer: C**

Explanation:

Feedback:

This targets the provisioning layer where inserted records and later attribute updates must both be captured and applied. When initial boarding records arrive but corrected passenger categories remain unchanged, the likely issue is update-event handling rather than calculation view activation.

## Question: 9

A regional courier company provisions parcel scan events into SAP HANA Cloud for same-day operations analytics. New scan records appear in the replicated target table, but several scans arrive without their related route assignment values. Source validation shows that route assignments are created shortly after the scan record and later updated in the source system.

The consuming calculation view activates successfully, and dashboard users can see the scan records but cannot group them correctly by route. The team must keep the existing target structure and avoid adding route correction logic in downstream models. The correction must identify why later route assignment changes are not reflected after the initial scan insert.

Which action best addresses the likely provisioning dependency?

Response:

- A. Add a calculation view expression that assigns ungrouped scan records to a default route during preview.
- B. Validate that the replication design captures route assignment updates after scan insertion and applies them to the existing target rows.
- C. Create a separate replicated table only for route assignments and ask each consuming model to join it independently.
- D. Manually update the missing route values in the target table before each same-day operations refresh.

**Answer: B**

Explanation:

Feedback:

This targets the provisioning layer where inserted scan records and later route assignment updates must both be captured and applied. When initial records arrive but subsequent assignment values remain missing, the likely issue is update handling after insert rather than calculation view activation.

## Question: 10

A construction materials company maintains SAP HANA Cloud modeling artifacts in SAP Business Application Studio. A developer updates a shared calculation view to expose a revised plant capacity attribute and activates the design-time object. During validation, the runtime artifact shows the revised attribute, but one dependent model still fails because its projection references the previous capacity field.

The environment follows controlled activation and deployment of development artifacts. The team must preserve lifecycle governance and avoid manual edits to deployed runtime objects. The correction must align dependent artifacts before the revised model is released to analytical consumers.

Which action best resolves the dependent model failure?

Response:

- A. Ask analytical consumers to refresh their metadata because the shared runtime artifact already exposes the revised capacity attribute.
- B. Manually edit the deployed runtime object so it exposes both previous and revised plant capacity fields.
- C. Update the dependent design-time projection reference and activate the affected artifacts in the correct deployment sequence.
- D. Create a duplicate shared calculation view with the previous capacity field so dependent models can remain unchanged.

**Answer: C**

Explanation:

Feedback:

This addresses the correct lifecycle and binding layer by aligning the dependent design-time reference with the revised shared model and deploying affected artifacts in sequence. When runtime is current but a dependent projection still references the old field, the correction belongs in dependent artifact alignment.

**Topic: 2**

**Unified Scenario Exam**

## Question: 11

### **CHALLENGE 1 — Recent Claim Status Access for Workload Review**

Velorian's claim status view performs acceptably for one agency region, but slows when all regions and recent claim status events are selected during workload review.

What should the data engineer validate first?

Response:

- A. Whether the calculation view consumes the intended virtual claim status event tables and how they behave under the all-region selection.
- B. Whether all claim status events can be replicated overnight so supervisors no longer depend on virtual event access.
- C. Whether claims supervisors can restrict the morning review to one agency region during the pilot.
- D. Whether workload formulas can be simplified by removing agency and claim type attributes.

**Answer: A**

Explanation:

Feedback:

The business requires recent claim status visibility, so the first validation point is the dependency between the calculation view and the intended virtual tables. Testing that dependency under the all-region selection addresses the performance exposure without changing the workload review requirement.

## Question: 12

### **CHALLENGE 1 — Recent Claim Status Access for Workload Review**

Claims supervisors need recent claim status activity, while agency reference attributes change infrequently.

Which approach best balances data currency and performance?

Response:

- A. Keep all claim status events and agency attributes virtual so every model follows a single access pattern.
- B. Keep freshness-critical claim status events virtual and validate whether stable agency attributes should remain staged locally.
- C. Replicate all claim status event and agency attribute data every night for predictable analytical performance.
- D. Remove agency attributes from workload review reporting to prevent joins with virtual claim status events.

**Answer: B**

Explanation:

Feedback:

This separates recent operational events from stable supporting attributes based on freshness need. It preserves current claim status visibility while reducing unnecessary remote workload from data that does not require virtual access.

## Question: 13

### **CHALLENGE 2 — Replicated Settlement History for Cycle Analysis**

Settlement-cycle reporting covers completed claims and does not require minute-level updates, but calculations run slowly when joined directly to virtual claim status access.

What is the most appropriate data provisioning choice?

Response:

- A. Use replicated completed settlement history in SAP HANA Cloud for cycle-time transformations and validate totals across agency selections.
- B. Keep completed settlement activity fully virtual so it matches the access pattern used for recent claim status events.
- C. Ask operations analysts to export virtual status data and calculate cycle-time trends outside SAP HANA Cloud.
- D. Remove adjuster group analysis from the cycle-time model to reduce query complexity.

**Answer: A**

Explanation:

Feedback:

Completed settlement history is performance-sensitive and does not require live currency, making replication appropriate for cycle-time trend calculations. Validating agency-level totals ensures the transformed data supports the reporting scope.

## Question: 14

### **CHALLENGE 2 — Replicated Settlement History for Cycle Analysis**

Cycle-time totals differ when analysts compare transformed settlement history with direct virtual claim status access.

Which dependency should be investigated before changing the final cycle-time measure?

Response:

- A. Whether the transformation logic, completed settlement record grain, and claim-to-status relationship are aligned before aggregation.
- B. Whether analysts can accept approximate cycle-time totals during the first controlled pilot phase.
- C. Whether all claim type attributes should be converted to virtual tables.
- D. Whether user display access differs between claims supervisors and operations analysts.

**Answer: A**

Explanation:

Feedback:

The mismatch appears between transformed settlement records and virtual status access, so the likely dependency is transformation grain and claim-to-status alignment. Changing the final measure before validating that relationship can hide the real cause.

**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/>**