

Boost up Your Certification Score

SAP C_BDCDA

**SAP Certified - Data Architect - SAP Business Data Cloud
(C_BDCDA_2605)**



For More Information – Visit link below:

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

Product Version

- ✓ Up to Date products, reliable and verified.
- ✓ Questions and Answers in PDF Format.

Visit us at: <https://www.examsboost.com/test/c-bdcda>

Latest Version: 4.0

1. Unified Scenario Simulation
2. Micro Skill Drill Exam

Topic: 1

Unified Scenario Simulation

Tadjoura Salt Works: One Governed Data Model on SAP BDC

Business context

Company

Tadjoura Salt Works is a solar sea-salt producer in Djibouti City, Djibouti, harvesting salt from shallow coastal evaporation pans and refining it for food-grade and industrial buyers across the Horn of Africa. Its work divides into a few clear areas: the evaporation pans, where seawater is drawn in and left to dry so salt can be scraped and gathered; a washing-and-refining plant that cleans and grades the raw salt; an iodization-and-quality lab that adds iodine to food-grade product and checks purity; and a bagging-and-dispatch team that packs and ships to wholesalers and factories under agreed contracts. Output swings with the dry season and with contract demand, and the firm has grown as more regional buyers have signed longer supply agreements.

Landscape

The works has recently begun using SAP Business Data Cloud to bring its harvest, refining, quality, and dispatch data into one governed place. For years its official month-end figures have been assembled by hand in an ageing shared spreadsheet workbook that each area feeds and everyone trusts. Leadership wants the data organised instead as governed, reusable data products that every area can consume from one place, and it wants that foundation trustworthy enough to support later AI-assisted analysis — without being torn up and rebuilt each time a new question is asked.

Team Formation

A small data team was formed to run the initiative. Its members are capable with reports but new to modelling data so it can be published as reusable products on a shared platform, and they are keen to show a result quickly. They have started sketching tables directly from the sales system, before pausing to agree what the underlying business concepts actually mean or how they relate to one another across the areas.

Modelling Request

The commercial lead asked for a sales summary showing tonnage sold by buyer and grade. One modeller began building platform-specific tables shaped only to that one summary, pulling fields straight from the sales system. Trouble surfaced almost at once: the pans, the refining plant, and dispatch each use the word "batch" to mean a different thing, and the sales system's idea of a "buyer" quietly lumps food-grade and industrial customers together in a way the quality lab and dispatch do not. A view shaped to that one report would not be a product any other area could reuse.

Two Layers of Modelling

During a review it became clear the team had jumped straight to physical tables without first agreeing the meaning of its core concepts or drawing how the entities relate. A platform advisor noted that a model meant to become a reusable data product is normally settled in order — first the agreed business meaning, then the entities and their relationships, then the platform-specific tables — and that skipping the earlier layers had baked one report's narrow shape into the design.

Governance Setup

At the same time, the team had begun bringing buyer-contract prices and iodization test results onto the platform with no record of which fields are sensitive or who owns them, and access was being

handed out one request at a time whenever someone asked. If the works is to treat this data as a dependable asset — including for the AI-assisted analysis leadership has in mind — sensitive figures need to be protected the same way for everyone, not guarded by whoever happens to field the request, so leadership wants fields classified and access governed by a consistent approach rather than ad-hoc permissions.

Legacy Reports

The old month-end workbook still runs, and some on the team want to switch it off the moment the first governed product appears, to be rid of a second source. Others point out that the areas have not yet checked the new figures against the workbook or come to rely on them, and that pulling the familiar source away too early would leave everyone without a trusted fallback while the new foundation is still proving itself.

Direction

The initiative lead wants the team to slow down just enough to agree what its core concepts mean before modelling them, to build the model in order from meaning to relationships to tables so the result can be reused as a governed data product rather than a single report's view, to classify sensitive fields so governance has something to act on, to protect access through a consistent framework instead of case-by-case grants, and to move off the ageing month-end workbook in phases rather than an abrupt cut-over. The reports themselves are not the hard part; the task is to establish a governed model and a governance approach the whole works can trust as a foundation for later analysis, without encoding one report's view, letting sensitive figures circulate unevenly, or stranding the areas as the old source is retired.

Your assignment

You are a data-architecture consultant advising the data team standing up SAP Business Data Cloud at Tadjoura Salt Works, a solar sea-salt producer. Advise the team so that its harvest, refining, quality, and dispatch data becomes a governed, reusable data product and a trustworthy foundation for later AI-assisted analysis, moving off an ageing hand-built month-end workbook. Five issues need your judgement: the areas use "batch" and "buyer" to mean different things and no shared meaning has been agreed; the modeller has begun building platform-specific tables before any conceptual or logical model; sensitive prices and test results are arriving unclassified; access is being granted one request at a time; and some want to switch off the legacy workbook the moment the first product appears. Recommend agreeing meaning before modelling, building the model in order from conceptual to logical to physical, classifying metadata so governance can act on it, governing access through a consistent framework, and phasing the move off the legacy source. Do not recommend encoding one report's view, granting access case by case, or an abrupt cut-over of the old source

1. CHALLENGE 1 — Agreeing What the Business Concepts Mean Before Modeling

- Advise the team to agree a shared meaning for its contested concepts, such as what a "batch" and a "buyer" are across the areas, before modelling them.
- Direct the team to record that agreed meaning as a conceptual model that the later work rests on.
- Explain why building around one area's meaning of a shared term leaves the governed figures untrustworthy for the others.

Checkpoints

- A written, agreed definition of the contested concepts exists before modelling begins.
- The areas consume the same combined figure the same way rather than by several different meanings.

2. CHALLENGE 2 — Ordering the Model From Conceptual to Physical

- Advise the team to build the model in order — agreed meaning, then entities and relationships, then platform-specific tables.
- Explain why jumping straight to physical tables shaped to one report encodes that report's view.

- Recommend that a logical model of entities and relationships precede any physical tables.

Checkpoints

- A logical model of entities and relationships exists before physical tables are built.
- The model can be reused beyond the single sales summary it started from.

3. CHALLENGE 3 — Classifying Metadata So Governance Can Be Applied

- Advise the team to classify each field by sensitivity, ownership, and business domain as it is brought in.
- Explain why unclassified sensitive fields leave governance nothing to act on.
- Recommend that classification precede any decision about who may access a field.

Checkpoints

- Sensitive fields are classified and their owners recorded before access is decided.
- Governance can point to a field's recorded sensitivity and owner when applying protection.

4. CHALLENGE 4 — Applying Governance as a Consistent Framework, Not Case-by-Case Permissions

- Advise the team to base access on a consistent framework driven by the recorded classifications.
- Explain why granting access one request at a time protects sensitive figures unevenly.
- Recommend how to meet real access needs without reverting to ad-hoc, per-request grants.

Checkpoints

- Access decisions follow the classification-based framework rather than one-off grants.
- The same level of sensitivity is protected the same way for everyone.

5. CHALLENGE 5 — Phasing the Move Off the Legacy Month-End Reports

- Advise the team to keep the ageing month-end workbook running alongside the governed product until the product is trusted and adopted.
- Explain why an immediate big-bang cut-over leaves the areas without a trusted fallback.
- Recommend retiring the legacy source only after the governed product is confirmed to agree with it and has been adopted.

Checkpoints

- The legacy workbook is retired only after the governed product is trusted and the two are confirmed to agree.
- No area is left without reliable month-end figures during the move.

Question: 1

CHALLENGE 1 — Agreeing What the Business Concepts Mean Before Modeling

The evaporation pans call a day's scraped salt a "batch," the refining plant calls a processed lot a "batch," and dispatch calls a shipment a "batch," and the modeller is about to shape the platform tables around only one area's meaning.

How should the advisor guide the team so the governed figures reflect a shared meaning every area can consume?

- Shape the tables around the refining plant's meaning of "batch" because that area is the busiest, and let the pans and dispatch rework their own reports to fit whatever the plant's tables end up representing.
- Let each area keep its own separate meaning of "batch" in its own tables, and leave whoever opens a combined report to remember which area's definition a given figure happens to be using that day.
- Agree a shared definition of each contested concept, such as what a "batch" means across the areas, and record it as a conceptual model before any platform tables are built.

D. Defer defining "batch" until the tables already exist, since the meaning will settle once the figures appear on screen and anything the areas still dispute can simply be renamed afterward without much trouble.

Answer: C

Explanation:

A governed data product is only reusable if every area reads its figures the same way, so the shared meaning has to be agreed and recorded as a conceptual model first. Settling it before any tables are built is what lets harvest, refining, and dispatch consume one trustworthy figure rather than three conflicting ones.

Question: 2

CHALLENGE 1 — Agreeing What the Business Concepts Mean Before Modeling

A modeller proposes agreeing a shared definition for "batch" but leaving "buyer" exactly as the sales system groups it — food-grade and industrial customers lumped together — and calls that good enough to start building.

Why is that only a partial fix rather than a complete conceptual model?

- A. Because "buyer" is contested too, so a conceptual model that settles only "batch" still leaves the combined figures untrustworthy for the quality lab and dispatch, which distinguish food-grade from industrial customers the sales grouping blends.
- B. Because the definition of "batch" itself should be written down separately for the pans and for dispatch, so the flaw is that a single agreed definition was recorded once rather than the sales system's grouping of "buyer" being left untouched in place across the areas.
- C. Because "buyer" is a simpler concept than "batch" and ought to have been defined first, so the only shortcoming here is the order in which the two shared terms were agreed rather than either being left unresolved.
- D. Because the sales system's grouping of "buyer" is already correct for every area of the works, so the model is in fact complete and describing it as a partial fix understates how much the team has genuinely resolved.

Answer: A

Explanation:

A conceptual model has to cover the concepts that are actually contested, and "buyer" is one of them, so settling only "batch" leaves the food-grade and industrial figures blended in a way other areas cannot use. The eventual data product must mean one thing to every area, which this partial definition does not deliver.

Question: 3

CHALLENGE 2 — Ordering the Model From Conceptual to Physical

Under pressure to show a sales summary quickly, the modeller has begun building platform-specific physical tables shaped to that one report before any model of the entities and how they relate.

Which sequence is complete, correctly ordered, and sound for a model meant to become a reusable data product?

- A. Build the platform-specific physical tables first, then draw the logical model of the entities and their relationships, then agree the conceptual meaning of the concepts, then let the areas consume the resulting data product.
- B. Agree the conceptual meaning first, then model the entities and their relationships as a logical model, then build the platform-specific physical tables, then let the areas consume the resulting data product.
- C. Model the entities and their relationships first, then build the physical tables, then agree what the concepts actually mean, then let the areas consume the resulting data product once the tables are already in place.
- D. Build the physical tables and let the areas start consuming them, then fit a logical model around whatever those tables happen to contain, and only then agree the conceptual meaning if disagreements later come to light.

Answer: B

Explanation:

Meaning comes first, then the entities and relationships that express it, and only then the platform tables that implement them. Each layer depends on the one before it, so this order yields a data product that reflects the business rather than a single report's view.

Question: 4

CHALLENGE 2 — Ordering the Model From Conceptual to Physical

The modeller wants to know the real risk of jumping straight to physical tables shaped to the one sales summary and skipping the logical model of entities and relationships.

What is the most likely risk?

- A. The tables simply refresh a little faster, since skipping the logical model removes a step, and the only real cost is that the team ends up writing slightly less documentation about how each of the figures is derived for the areas that consume them.
- B. The tables become easier to reuse, because a physical design shaped to one report already captures every relationship the other areas could ever need without any further modelling of entities being required at all.
- C. The physical tables encode one report's shape and never capture how harvest, refining, and shipment relate, so the result serves the single summary but cannot be reused as a data product by any area whose questions it did not cover.
- D. There is no real risk at all, since a physical model built for one report can always represent every entity and relationship in the whole business simply by adding a few more columns to the same tables later on.

Answer: C

Explanation:

Skipping the logical model means the relationships among harvest, refining, and shipment are never captured, so the tables can serve the one summary but cannot become a reusable data product, breaking down as soon as another area asks something the report did not.

Question: 5

CHALLENGE 3 — Classifying Metadata So Governance Can Be Applied

Buyer-contract prices and iodization test results are being brought onto the platform with no record of which fields are sensitive or who owns them.

What should the team do first so that governance can later be applied to the data product?

- A. Classify each field by sensitivity, ownership, and business domain as it is brought in, so that later access and protection decisions have those recorded classifications to rest on.
- B. Bring in every field with the same open setting first and wait to see which of them cause complaints, then classify only those particular fields once someone actually objects to how the data is being used.
- C. Leave the fields unclassified and rely on each analyst to remember which figures are sensitive, since the people working with the data know its sensitivity far better than any recorded classification ever could.
- D. Restrict every field to the modelling team alone for now so that nothing needs classifying, and simply withhold access from everyone else until the entire model is eventually finished and then formally reviewed.

Answer: A

Explanation:

Classifying fields by sensitivity, ownership, and domain as they arrive is the foundation governance later acts on. Without those classifications there is nothing for access rules to reference, so recording them first is what lets the data product be protected consistently.

Question: 6

CHALLENGE 3 — Classifying Metadata So Governance Can Be Applied

A set of sensitive contract prices was seen by someone who should not have had them, even though everyone believed the data on the platform was being handled carefully.

What is the most likely underlying cause?

- A. The platform simply failed to protect the field on its own, so the sensible response is to expect the platform to guard sensitive data automatically rather than to record anything about the field beforehand.
- B. The owner of the data must have approved the access in error, so the cause is one mistaken human decision rather than anything about how the field itself was described or classified on the way onto the platform.
- C. The report refreshed too quickly for the access check to complete, so the exposure is really a timing problem in how the figures were rendered rather than any gap in how the field's sensitivity had been recorded on the way onto the platform.

D. The field was never classified by sensitivity or owner, so the access rules had nothing to enforce and a price that should have been restricted travelled through the platform as if it were ordinary, unremarkable data.

Answer: D

Explanation:

Access rules can only act on a field's recorded sensitivity and owner, so an unclassified price has nothing marking it as restricted and moves like ordinary data. The missing classification is the underlying cause, which is why classifying fields as they arrive is what prevents this.

Question: 7

CHALLENGE 4 — Applying Governance as a Consistent Framework, Not Case-by-Case Permissions

Access to the sensitive contract prices is currently being granted one request at a time, whenever an analyst happens to ask a colleague for it.

What is the best recommendation?

- A. Keep granting access one request at a time as people ask, since handling each case on its own keeps the team flexible and there is no need for a shared rule as long as the individual requests are all being answered.
- B. Grant everyone the same broad access to the sensitive fields to save time, since a single open setting is simpler to administer than either per-request grants or a classification-driven governance framework would be.
- C. Base access on a consistent framework driven by the recorded classifications, so the same sensitivity is protected the same way for everyone rather than depending on who asked whom for a given field.
- D. Route every access request to a single senior manager to approve in person, so the decisions stay case by case but at least pass through one individual rather than resting on any recorded classification of the fields.

Answer: C

Explanation:

A classification-driven framework applies one consistent rule to each level of sensitivity, so protection no longer depends on individual requests. That evenness is what makes the governed data product dependable for everyone the same way.

Question: 8

CHALLENGE 4 — Applying Governance as a Consistent Framework, Not Case-by-Case Permissions

The initiative lead, who wants the governed data ready to support later AI-assisted analysis, asks what the works actually gains by protecting access through a consistent framework rather than case by case. Which benefit best supports that choice?

- A. A consistent framework protects the same sensitivity the same way for everyone, so the governed data stays a dependable, trustworthy foundation that other areas — and later AI-assisted analysis — can rely on rather than one guarded differently each time.
- B. The main gain is that the team can grant access more quickly to whoever asks, since a standing framework speeds up provisioning, so its value lies in faster turnaround on individual requests rather than in any evenness of how the data is protected.
- C. A framework matters only once the data is actually feeding an AI model, so until that point granting access one request at a time is equally trustworthy and there is little to be gained by adopting the framework any earlier than the analysis truly begins.
- D. The chief benefit is that fewer people need to be involved in each access decision, so the value lies mainly in the reduced administrative workload rather than in protecting the same sensitivity evenly for everyone who is entitled to it.

Answer: A

Explanation:

Treating data as a dependable strategic asset, including for later AI-assisted analysis, rests on it being governed consistently; a framework that protects each level of sensitivity evenly is what keeps the foundation trustworthy, which case-by-case grants cannot promise.

Question: 9

CHALLENGE 5 — Phasing the Move Off the Legacy Month-End Reports

The works still assembles its official month-end figures in an ageing shared spreadsheet workbook. The team wants to switch every area off it the moment the first governed data product is ready.

What is the best approach to moving off the legacy workbook?

- A. Switch every area off the old workbook the day the first governed product exists, since running two sources at once is wasteful and a clean immediate cut-over is the simplest way to force the areas to adopt the new foundation.
- B. Keep the old workbook running alongside the governed product until the product is trusted and adopted, then retire it, so the move off the legacy source is phased rather than a single abrupt cut-over.
- C. Keep the old workbook as the official month-end source indefinitely and treat the governed product as a secondary view, since the areas already trust the spreadsheet and there is little to gain from ever retiring it.
- D. Abandon the governed product and simply tidy up the old workbook instead, since modernizing the month-end process adds risk that a familiar, already-trusted spreadsheet does not carry for the areas at this stage.

Answer: B

Explanation:

Modernizing a familiar source is safer done in phases than as a big-bang cut-over, so running the old workbook alongside the governed product preserves a trusted fallback until the areas have confirmed the new figures and adopted them, and only then is the legacy source retired.

Question: 10

CHALLENGE 5 — Phasing the Move Off the Legacy Month-End Reports

A colleague pushes to retire the month-end workbook immediately, before the governed product has been checked against it or adopted by the areas.

Why is that immediate cut-over risky?

- A. It is not really risky, because a governed product is correct as soon as it is built, so there is never any need to run it alongside the old workbook or to confirm that the two agree before the spreadsheet is retired for good.
- B. The only risk is a little duplicated effort, since running the workbook a while longer merely wastes some time but has no real bearing on whether the areas trust the governed figures once the old source has been removed.
- C. The risk is purely the technical storage cost, since keeping the workbook consumes space, so retiring it early is the prudent, lower-cost choice regardless of whether the areas have actually adopted the governed product yet.
- D. Cutting the old source off before the governed product is checked and adopted leaves the areas with no trusted fallback and forces reliance on figures they have not yet confirmed, which a phased coexistence would avoid.

Answer: D

Explanation:

Retiring the workbook before the governed product is confirmed to agree and has been adopted removes the areas' trusted fallback and makes them depend on unproven figures. Phased coexistence keeps the old source available until confidence is established, which is why an immediate cut-over is the risk.

Topic: 2

Micro Skill Drill Exam

Question: 11

Aurelia Assurance, a property-and-casualty insurance group in Dublin, Ireland, is starting an analytics initiative on SAP Business Data Cloud to support claims reporting and, later, AI-assisted underwriting decisions. Under schedule pressure, a project team has begun building physical tables directly in the target platform, naming columns after the screens and report layouts staff already use, before the business has agreed what the core entities — policy, claim, and coverage — actually mean. Business analysts and claims leads still disagree on some of those definitions. The lead data architect is asked how the model should be established so that it reflects agreed business meaning and stays stable as reports evolve and new AI use cases appear. The constraint is explicit: the model must serve many downstream reports and future analytical workloads, not just today's screens, and reworking the foundation later would be costly in a regulated environment. The architect must choose the approach that fixes shared meaning before platform-specific structure is committed.

How should the architect establish the data model?

- A. Keep building the physical tables directly in the target platform to mirror the existing report layouts, so the team preserves momentum and delivers the first reports on the date already promised to the business.
- B. Have each downstream report owner define its own tables independently, so every report gets exactly the structure it needs without waiting for shared entity definitions to be agreed across teams.
- C. Reuse the entity definitions copied from a similar past project at another regional office, so the model rests on already-proven structures instead of re-deriving the business concepts from scratch here.
- D. Define a conceptual and logical model of the agreed business entities first, then derive the physical structures from it, so the design reflects shared business meaning and stays stable as reports change.

Answer: D

Explanation:

Defining a conceptual and logical model of the agreed business entities and their relationships fixes the shared meaning before any platform-specific structure exists, and deriving the physical model from it keeps the design aligned to those concepts and stable as reports and later AI needs evolve. This acts at the foundational modeling layer rather than the presentation layer.

Question: 12

Cordillera Beverages, a bottling and beverage manufacturer in Bogota, Colombia, is consolidating production, distribution, and loyalty data onto SAP Business Data Cloud to support a widening range of analytics and AI use cases across sales and supply. The current legacy approach transforms and pre-aggregates all data in a separate staging tool before loading only pre-shaped summary tables, so every new business question forces the upstream transformation to be re-engineered before it can be answered. Data volumes are large and growing, and the target platform is capable of transforming data at scale on its own. The architect must recommend an integration pattern that keeps the raw detail available and lets new questions be served without rebuilding upstream jobs each time. The constraint is that analytics needs are still emerging and will keep changing, so the design cannot assume the questions are all known in advance. Momentum matters, but so does avoiding another rigid pipeline that has to be re-engineered for every new question.

Which integration pattern should the architect recommend?

- A. Keep the existing ETL approach that transforms and pre-aggregates all data before loading, so only clean, ready-to-use summary tables land in the platform and downstream users never handle raw detail at all.
- B. Adopt an ELT pattern that loads the raw detail into the capable target platform first and transforms it there per use case, so new questions are served without re-engineering upstream jobs.
- C. Replace the pipeline with periodic manual extracts that analysts shape in spreadsheets, so each team can transform the data in the tool it already knows without depending on a central platform.
- D. Add more pre-built ETL jobs for every anticipated report, so a ready transformed table already exists in advance for each question the business is expected to ask this year.

Answer: B

Explanation:

Loading the raw detail first and transforming it inside the capable target platform keeps full granularity available and lets each new use case define its own transformation without rebuilding upstream jobs. This fits large, growing volumes on a platform that can transform at scale and matches questions that are not all known in advance.

Question: 13

Talavera Drivetrain, an automotive parts manufacturer in Guadalajara, Mexico, is unifying data from HR, warranty, and supplier systems onto SAP Business Data Cloud. As each team requests access to particular datasets, administrators have been granting the requests one by one, and over time some sensitive personal and commercial fields have become widely accessible. An internal audit now flags inconsistent handling of sensitive data and warns that the problem will grow as more sources are onboarded. The architect is asked to fix access in a way that will scale rather than be repaired dataset by dataset. The constraint is clear: sensitivity must be handled consistently across all current and future datasets, driven by what the data is rather than by who happens to ask for it. Leadership wants a defensible, repeatable basis for access that an auditor can follow, not another round of individual approvals. The architect must decide where the real fix belongs rather than treat only the fields this audit named.

What should the architect recommend to fix access so it scales?

- A. Keep granting each access request individually but add a manager approval step, so every grant is signed off while the request-by-request process the teams already know is otherwise preserved unchanged.
- B. Restrict all sensitive datasets to administrators only and have them run reports on behalf of other teams, so exposure falls because fewer people hold direct access to the data.
- C. Establish a data governance framework that classifies metadata by sensitivity and drives access from those classifications, so every current and future dataset is handled consistently by policy.
- D. Mask the specific fields the audit named across the affected datasets, so the exposed sensitive values are hidden and the immediate finding from this particular audit is resolved for those tables.

Answer: C

Explanation:

Classifying metadata by sensitivity and driving access from those classifications within a governance framework applies one consistent policy to every dataset, present and future, so access follows the data's classification rather than individual requests. This acts at the governance layer where the audit gap actually sits.

Question: 14

Meridian Trust Bank, a retail bank in Singapore, is defining how a new enterprise data architecture on SAP Business Data Cloud aligns with the company's broader transformation programme. Leadership has adopted TOGAF as the enterprise-architecture framework, while the data office works to DAMA-DMBOK for its data-management disciplines — governance, data quality, and metadata. A steering group, worried about duplicated effort, proposes standardizing on a single framework and asks the architect to recommend keeping TOGAF, keeping DAMA-DMBOK, or choosing just one of them. The constraint is

that the resulting design must both align to the enterprise architecture leadership has committed to and fully cover the data-management disciplines the data office is accountable for. The architect knows the two frameworks were built for different concerns, and that dropping either one would leave a real gap. She must advise the steering group on how to treat the two frameworks without leaving the design half-covered.

How should the architect advise the steering group on the two frameworks?

- A. Use TOGAF and DAMA-DMBOK together as complementary frameworks — TOGAF for enterprise-architecture alignment and DAMA-DMBOK for data-management disciplines — so both the enterprise fit and the data disciplines are covered.
- B. Standardize on TOGAF alone and drop DAMA-DMBOK, so the data architecture stays tightly aligned to the enterprise architecture leadership has already committed to, without a second overlapping framework to maintain.
- C. Standardize on DAMA-DMBOK alone and drop TOGAF, so the data-management disciplines are covered in depth while the team avoids the overhead of operating a separate enterprise-architecture framework alongside it.
- D. Defer both frameworks and let each project pick its own approach, so teams can move quickly now without being constrained by a framework decision the steering group has not yet been able to settle.

Answer: A

Explanation:

TOGAF and DAMA-DMBOK address different concerns — enterprise-architecture alignment and data-management disciplines — so using them together satisfies both requirements the design must meet. Treating them as complementary rather than redundant is what lets the architecture fit the enterprise while still covering governance, quality, and metadata.

Question: 15

Kalahari Power, an electricity utility in Windhoek, Namibia, federates live access to a large external market-and-grid operational platform alongside governed data already held in SAP Business Data Cloud, so analysts can combine both without copying everything in. Recently the dashboards that join one very large, frequently-queried external history table have become slow, and analysts are complaining. The team is debating how to fix the performance without losing what federation gave them. One faction wants to abandon federation and bulk-copy every external table in; another wants to keep federating everything and simply accept the latency. The architect must recommend how to resolve the slow dashboards while keeping data governed and avoiding needless duplication. The constraint is that only the one large, hot dataset is causing the problem, while the remaining federated tables perform acceptably and do not need copying. The architect has to target the fix to where the bottleneck actually is rather than change the whole landscape.

How should the architect resolve the federation performance problem?

- A. Keep federating every external table live and accept the slower dashboards, so nothing is duplicated and the governed single source is preserved exactly as it stands today.
- B. Abandon federation and bulk-copy all external tables into the platform on a schedule, so every query runs on local data and the latency disappears across every dashboard at once.

- C. Continue federating the external data but selectively replicate only the large, frequently-queried dataset, so the hot data performs while the rest stays federated and copies do not proliferate.
- D. Move the slow dashboards off the platform onto the external system entirely, so the heavy joins run where that data lives and the governed platform is left out of those cross-platform queries altogether.

Answer: C

Explanation:

Federating by default while selectively replicating only the large, hot dataset targets the problem exactly where it occurs, so that dataset performs locally while everything else stays federated and governed and copies do not proliferate. It matches the fix to the specific bottleneck rather than the whole landscape.

Question: 16

Astra Media, a digital media and streaming company in Reykjavik, Iceland, wants audience-engagement data to be reusable across domains and ready for AI-assisted analytics on SAP Business Data Cloud, instead of every team rebuilding its own private extract. The lead architect proposes publishing the data once as a governed data product that other domains can discover and reuse. During design review the team pushes back with a practical

question: what actually makes it a governed, reusable data product rather than just another dataset copy given a friendly name? The constraint is that the result must be discoverable and reusable by other domains, governed for sensitivity and quality, and must not simply be a private extract that consumers each reshape on their own.

The architect has to specify the complete, minimal set of elements the asset needs — enough to make it a true data product, but without adding unnecessary machinery that would recreate the copy-proliferation the product is meant to replace.

Which minimal set of elements should the architect require to publish it as a governed, reusable data product?

- A. Assign a clear owner, publish a raw table copy that each domain may reshape freely, register it for discovery, and apply governance and quality rules.
- B. Assign a clear owner, define a stable interface and access contract, register it for discovery with classified metadata, and apply governance and quality rules.
- C. Assign a clear owner, define a stable interface and access contract, register it for discovery with classified metadata, and stand up a separate duplicate warehouse holding a private copy for every consuming domain.
- D. Assign a clear owner, define a stable interface and access contract, announce it and take change requests by email, and apply governance and quality rules.

Answer: B

Explanation:

An owner makes it accountable, a stable interface and access contract makes it dependable to reuse, discovery with classified metadata makes it findable and governed for sensitivity, and governance and quality rules keep it trustworthy — the minimal complete set that turns a dataset into a governed, reusable data product, with nothing required missing and nothing unnecessary added.

Thank You for Trying Our Product

For More Information – **Visit link below:**

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

15 USD Discount Coupon Code:

G74JA8UF

FEATURES

- ✓ **90 Days Free Updates**
- ✓ **Money Back Pass Guarantee**
- ✓ **Instant Download or Email Attachment**
- ✓ **24/7 Live Chat Support**
- ✓ **PDF file could be used at any Platform**
- ✓ **50,000 Happy Customer**



Visit us at: <https://www.examsboost.com/test/c-bdcda>