

SAP

P_BTPA_2511

SAP Certified - Solution Architect - SAP BTP

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/p-btpa-2511>

Latest Version: 4.1

Subjects

1. Regional Transport Authority BTP Data Discovery Advisory
2. Cedarline Health Network's First BTP Integration Pilot
3. Meridian Freight Group BTP Architecture Kickoff Alignment
4. Specialty Retailer's First SAP BTP Account Foundation
5. Aldermont Energy Trading BTP Extension Delivery Foundation
6. Consumer Goods Group BTP Rollout Cost and Consumption Plan
7. Meridian Assurance BTP Extension Design for Underwriting
8. Regional Water Utility BTP Data and Analytics Blueprint
9. Telecom Operator Intelligent-Enterprise Blueprint on SAP BTP
10. Haldenmark Automotive Systems Scaling Its BTP Landscape
11. Nordhavn Biologics — Governing Embedded AI Across a Regulated BTP Landscape
12. Global Manufacturer BTP Integration Architecture Review
13. Classifying integrations and opening partner connectivity on a first platform
14. Standing up a first platform footprint and sizing a first commitment
15. One governed source of truth for bottling analytics
16. A side-by-side extension and governed automation on SAP BTP
17. Laying a governed SAP BTP account foundation for a cocoa refiner
18. Setting the architecture foundation for a new platform programme
19. Anchoring a target landscape in reference architecture and governing its build
20. Consolidating integration and grounding intelligence on a governed foundation
21. Grounding generative assistance and sizing a seasonal platform investment
22. Rationalizing fragmented integration across a regulated freight landscape
23. Turning strategy into a governed, traceable, affordable platform blueprint
24. Placing a shared AI foundation and governing agents across regulated regions
25. Integration Architecture
26. Data & Analytics Architecture
27. AI Services Architecture
28. Application Development, Automation & Extension Methodology
29. BTP Foundation Services & Solution Governance
30. Enterprise Architecture Roles & SAP EA Framework
31. BTP Commercial Models & Cost Estimation
32. SAP BTP Foundations, Account Model & Environments

Topic: 1

Regional Transport Authority BTP Data Discovery Advisory

Business Context

Agency Context

The Regional Transport Authority is a mid-sized public-sector body that plans and funds passenger transport across a metropolitan region. It has recently begun a data initiative and has engaged an SAP BTP solution architect for an early discovery engagement. The authority runs SAP applications for finance and asset management, alongside several non-SAP operational systems for ticketing and scheduling. Leadership wants to make better funding and service decisions from its own data, but the

organization has little prior experience treating data as a shared, reusable asset rather than as system-specific extracts.

Discovery Conversations

Across the first discovery workshops, a recurring pattern emerges. Whenever a department needs figures, an analyst asks the finance or asset teams for a fresh extract, and each request is fulfilled as a one-off file shaped to that request. The planning team, the grants team, and the internal audit function each hold their own copy of what is essentially the same ridership-and-cost information, pulled at different times with slightly different filters. Nobody can say confidently which copy is authoritative, and the same monthly figure often appears with small differences in two reports presented to the board.

Consumption Landscape

Several downstream uses for the data are already lining up. A public transparency dashboard is planned, a grants-reporting pack must be produced for a national funding body, and the planning team wants to explore analytics on service coverage. Each of these consumers has been scoped, so far, as its own pipeline pulling directly from the source systems. The architect notes that the ridership and subsidy figures at the center of all three uses are the same underlying information, described and secured differently by each team.

Governance Context

Because the authority is a public body, some of the data carries residency and disclosure obligations, and passenger-level records are more sensitive than the aggregated figures used for public reporting. The transparency dashboard is intended for open publication, while internal audit needs fuller detail. Today those boundaries are handled informally, decided per extract by whoever prepares the file. There is interest in SAP Business Data Cloud as the direction for the authority's data foundation, though the team is still learning what a governed data foundation would mean for how they work.

Delivery Pressure

The grants-reporting deadline is the nearest fixed date, and one manager has suggested simply scripting a direct pull from the finance system into a spreadsheet for that report, to be safe. The planning lead, meanwhile, is keen to start the coverage analytics quickly and has asked whether her team can just take a copy of the raw source tables for now and sort out governance later. Both requests would meet an immediate need, and both would add another uncontrolled copy of the same figures to a landscape that already struggles to reconcile them.

Advisory Focus

The sponsor has asked the architect for a clear early recommendation on how the authority should position this shared ridership-and-cost information for the several consumers now emerging, before any of the individual pipelines are built. The engagement is still at discovery, so no target build has been committed; the value of this stage is in setting the right foundation direction rather than in delivering a specific report.

Your assignment

You are acting as the SAP BTP solution architect engaged by the Regional Transport Authority for an early data-discovery engagement. The authority is a public-sector body beginning its first structured data initiative. Several consumers of the same ridership-and-cost information are emerging at once, delivery pressure from a grants deadline is present, and the sponsor wants an early recommendation on how to position this shared data before individual pipelines are built. Reason as an advisor applying the Data and Analytics Advisory Methodology at concept level; do not assume a target build has been approved

1. CHALLENGE 1 — Positioning Shared Transport Data as a Governed Product

Advise the sponsor on whether the shared ridership-and-cost information should be positioned as a governed, reusable data product or delivered as separate extracts for each consumer.

Frame the recommendation in terms of reuse across the transparency dashboard, grants pack, and coverage analytics, at a discovery altitude, without committing to specific services or a build design. Explain what defining the data once as a product changes for consumers that appear later, including how ownership of that product should be understood.

Checkpoints

The recommendation results in one agreed definition of the shared data that all known consumers can reuse.

The advice remains at discovery altitude and does not depend on a specific build being approved.

2. CHALLENGE 2 — Reconciling Conflicting Copies of Ridership Figures

Advise where the divergent board figures should be addressed, distinguishing the visible reporting symptom from its upstream cause.

Position the fix so that future consumers, and short-term deadline workarounds, do not reintroduce the same divergence.

Checkpoints

The same monthly figure resolves identically across reports because it derives from one defined source.

The chosen action addresses the shared-source cause rather than only the report currently in view.

3. CHALLENGE 3 — Setting Data Boundaries for Public and Internal Use

Advise how the authority should handle the difference between openly publishable aggregated data and sensitive passenger-level detail for internal audit, given residency and disclosure obligations.

Position the boundary decision relative to the grants deadline pressure and the request to defer governance until later.

Checkpoints

Open and internal data are separated by a boundary defined on the shared source, not by per-extract filtering.

The recommendation weighs the nearest deadline against the cost of adding another ungoverned copy.

Question: 1

CHALLENGE 1 — Positioning Shared Transport Data as a Governed Product

The transparency dashboard, grants pack, and coverage analytics all rely on the same ridership-and-cost information, and each has been scoped as its own pipeline pulling from the source systems. The sponsor asks how this shared information should be positioned at discovery.

What is the best advisory recommendation?

- A. Let each consumer keep building its own extract now, then merge the separate definitions into one shared definition after all three pipelines are live.
- B. Define the shared ridership-and-cost information once as a governed, reusable data product that both the known and future consumers draw from.
- C. Give the first consumer to go live ownership of the data, and have the other teams request working copies from that team as they need them.
- D. Defer any positioning decision until the target data platform has been selected and its detailed build scope has been formally approved.

Answer: B

Explanation:

Positioning the shared information once as a governed, reusable data product means every consumer reuses one authoritative definition instead of duplicating it. This is the discovery-altitude foundation move, and it is exactly what the Data and Analytics Advisory Methodology structures decisions around — a governed data product for reuse rather than point pulls.

Question: 2

CHALLENGE 1 — Positioning Shared Transport Data as a Governed Product

A new consumer — a fare-policy review team — is expected to need the same ridership-and-cost information a few months after the first three consumers are set up.

How should the discovery recommendation account for consumers that appear later?

- A. Position the data as a reusable product now, so that a later consumer draws from the same governed definition without a new extract being built for it.
- B. Record the fare-policy team as out of scope for this particular discovery, and let that team commission its own separate pipeline later, when it eventually arrives.
- C. Ask each of the current consumers to retain a spare copy of its own extract, so that the fare-policy team can later reuse one of those copies.
- D. Wait to define the shared data until the fare-policy requirement is fully confirmed, so that a single definition can cover every consumer at once.

Answer: A

Explanation:

Defining the data once as a reusable product is precisely what lets a later consumer plug into the same governed definition rather than triggering yet another extract. Reuse is the core benefit of positioning data as a product at discovery, and it is what protects the authority as more consumers emerge.

Question: 3

CHALLENGE 1 — Positioning Shared Transport Data as a Governed Product

The sponsor accepts positioning the shared information as a data product but asks who, in practice, should be accountable for that product once it exists.

Which advisory position best fits a governed data product at discovery?

- A. Whichever team happens to consume the data most heavily should hold accountability, since that team has the strongest day-to-day interest in it.
- B. Accountability should rotate between the consuming teams each reporting period, so that no single team carries the ongoing effort of maintaining it.
- C. A defined steward should be accountable for the product's definition and its data boundary, so that all consumers rely on one governed, maintained source.
- D. No standing accountability is needed at discovery, since the product can simply be maintained informally by whoever last updated the underlying figures.

Answer: C

Explanation:

A governed data product needs a defined steward accountable for its definition and boundary, so every consumer relies on one maintained source rather than an unowned copy. Stewardship is the mechanism that keeps the shared definition authoritative as consumers change over time.

Question: 4

CHALLENGE 2 — Reconciling Conflicting Copies of Ridership Figures

The board has received two reports in which the same monthly ridership figure differs slightly. Each report was built by a different team from its own extract, each taken at a different time.

Where should the architect advise addressing the discrepancy?

- A. In the report that was presented most recently, by correcting its figure so that it matches the figure shown in the earlier report.
- B. By standardizing the filters that each team applies, so that their separate extracts are configured to produce the same monthly number.
- C. At the shared data source, by having both reports derive the figure from one governed definition rather than from separate extracts.
- D. In each of the two reports, by adding a reconciliation note that explains to the board why the two figures differ from each other.

Answer: C

Explanation:

The divergence is a symptom; its cause is that each team pulls its own copy at a different time. Resolving it at the shared source, so both reports derive the figure from one governed definition, removes the cause and keeps future reports consistent. This is the second-order fix rather than a per-report patch.

Question: 5

CHALLENGE 2 — Reconciling Conflicting Copies of Ridership Figures

A manager proposes to meet the near grants deadline by scripting a direct pull from the finance system into a spreadsheet for that one report only.

How should the architect respond at discovery?

- A. Approve the direct pull for the grants pack alone, treating it as an exception, since it is only a single report and it has a firm, non-negotiable submission deadline coming up.
- B. Advise against another direct one-off pull, and recommend that the grants pack draw from the shared governed definition being established for all consumers.
- C. Approve the direct pull for now, but ask the manager to delete the working spreadsheet as soon as the grants pack has been submitted.
- D. Redirect the grants pack so that it reuses the planning team's existing extract, rather than pulling a fresh copy from the finance system.

Answer: B

Explanation:

A direct one-off pull adds another uncontrolled copy of the same figures to a landscape that already cannot reconcile them, which is the exact pattern under review. Pointing the grants pack at the shared governed definition meets the deadline while reinforcing the foundation rather than undermining it.

Question: 6

CHALLENGE 2 — Reconciling Conflicting Copies of Ridership Figures

An analyst argues that the figures differ only because of a formula error in one dashboard, and that fixing that formula will make the two board figures agree.

How should the architect evaluate this explanation?

- A. The explanation is incomplete, because the figures come from separate extracts taken at different times, so aligning one formula leaves the divergence able to return.
- B. Accept the analyst's explanation at face value and have the formula corrected, since a matching calculation in both of the dashboards is ultimately what the board actually needs to see.
- C. Treat the formula as irrelevant and instead ask both teams to round their figures the same way, so that the small differences disappear from view.
- D. Defer the question to the board itself, since it is best placed to decide which of the two ridership figures it would prefer to treat as correct.

Answer: A

Explanation:

The figures differ because each report is built from its own copy pulled at a different time, so a formula fix would not address the shared-source cause. Recognizing the extract divergence as the real cause is the second-order reasoning the scenario requires, and it points back to one governed definition.

Question: 7

CHALLENGE 3 — Setting Data Boundaries for Public and Internal Use

The transparency dashboard will publish aggregated figures openly, while internal audit needs fuller passenger-level detail. Both draw on the same underlying records, and today sensitivity is decided per extract by whoever prepares the file.

What should the architect recommend?

- A. Keep deciding sensitivity per extract, but issue a written guideline that describes how preparers should filter each file before it is released.
- B. Build a separate, unrelated dataset for the public dashboard, so that its published data never touches the same underlying internal records at all.
- C. Publish the full detail to the dashboard, and rely on the dashboard's own access settings to hide the sensitive fields from public viewers.
- D. Define the data boundary on the shared source, so that aggregated open data and sensitive detail are separated by design for every consumer.

Answer: D

Explanation:

Defining the boundary on the shared source means the split between openly publishable aggregates and sensitive detail is enforced once, by design, for even, 'consumer — meeting residency and disclosure obligations consistently. It replaces the ad hoc per-extract decisions that risk leaking passenger-level detail into open publication.

Question: 8

CHALLENGE 3 — Setting Data Boundaries for Public and Internal Use

The planning lead asks to take a copy of the raw source tables now to start coverage analytics quickly, with governance to be sorted out later, while the grants deadline is close. How should the architect weigh this request at discovery?

- A. Allow the raw copy for the planning team now, on the understanding that the required data boundary will be retrofitted onto it before anything is published.
- B. Allow the raw copy but restrict it to the aggregated tables only, so that no passenger-level detail is included in the copy the planning team takes.
- C. Advise against an ungoverned raw copy, and recommend the analytics draw from the governed product with its data boundary applied, even under deadline pressure.
- D. Pause the coverage analytics work entirely until the full governed platform has been built and its boundaries are enforced end to end across the landscape.

Answer: C

Explanation:

Taking a raw copy "to sort out later" adds another ungoverned duplicate carrying sensitive detail, which is the risk the engagement is meant to remove. Recommending the analytics draw from the governed product with its boundary applied protects residency and disclosure obligations while still letting the work proceed.

Question: 9

CHALLENGE 3 — Setting Data Boundaries for Public and Internal Use

Legal advisers confirm that some passenger-level records carry residency and disclosure obligations, and they ask how those obligations should shape the shared data product. What is the best advisory position?

- A. Treat residency and disclosure as constraints on the shared product itself, so its boundary keeps obligated detail out of open uses while allowing audit its fuller view.
- B. Treat the obligations as the transparency dashboard's responsibility alone, since that dashboard is the only consumer that actually publishes anything to the public.
- C. Collect all records, including the obligated passenger-level detail, into one open pool, and address the obligations later once the dashboard's audience is better understood.
- D. Exclude the obligated records from the data product completely, so that neither the public dashboard nor the internal audit team can ever make use of them.

Answer: A

Explanation:

Residency and disclosure obligations belong on the shared product, so its boundary keeps obligated detail out of open uses while still serving audit's fuller view. Making obligations a property of the governed product is what applies them consistently to every consumer rather than one at a time.

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/>