SAP
C_BAIPA_2610
SAP Certified - Solution Architect - SAP Business AI Platform

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

Topic: 1
Unified Scenario Simulation
Kantele Rights Collective AI-Ready Royalty Platform Blueprint
Business Context
Kantele Rights Collective is a Finnish non-profit music-rights society headquartered in Tampere. It licenses broadcast, live-venue, and streaming use of music on behalf of roughly 38,000 composer, lyricist, and publisher members, collects the licence fees, and distributes royalties to members four times a year. Member services handles about 2,400 enquiries a month, most of them asking why a particular work earned less than expected. A queue of disputed and unidentified works has grown over the last two years and now delays part of every quarterly distribution run.
System Landscape
Finance runs on an SAP cloud ERP system that went live two years ago with a deliberately standard configuration, and the board adopted a clean-core principle at that go-live. Despite that intent, a former implementation partner left an undocumented custom view in the core, which an extension team's early dispute prototype now reads. Royalty calculation still happens in a distribution engine built in-house in the early 2000s, which holds rights-holder registrations and the split percentages agreed between co-writers. Usage reports from broadcasters arrive through a mix of file transfers and emailed spreadsheets, and a set of scheduled scripts reformats them for the distribution engine. An agency-built member portal lets members view statements and register new works, and several streaming services are due to start reporting usage to Kantele next year.
Programme Direction
The IT director, Aino Lehtonen, sponsors a design-validation phase on SAP Business AI Platform. Two outcomes are in scope: a Joule-based assistant that answers member payout questions, and an AI agent that triages unidentified-work disputes by proposing the most likely rights holder for reviewers. The design team has reviewed Joule Studio and SAP AI Agent Hub as part of the agent workstream. A contractor has already run a small agent pilot against three months of historical disputes, using a shared environment that also hosts developer experiments.
Pilot Observations
The pilot produced sensible proposals for most broadcasters. For two regional broadcasters, however, the agent suggested the wrong rights holder in roughly one case in six. Those two broadcasters still email spreadsheets, and the script that reformats their files was written by an employee who left last year. The contractor wants to retune the agent's instructions first, and has also offered to build three new point-to-point connectors, one per broadcaster file format, within two weeks.
Stakeholder Positions
Mikko Rautio, head of member services, wants the assistant live before the next quarterly distribution and would like it to explain co-writer splits in full. The data protection lead, Saara Heikkilä, notes that each member's individual payout amount is that member's personal financial data. The finance director, Jukka Peltonen, wants dispute status visible to finance, and the finance core lead has proposed changing standard ERP objects, adding dispute fields to the standard open-item records and coding an approval step into the standard payment run, which could go live this quarter. Dispute records themselves currently live in the distribution engine and in member-services spreadsheets. The contractor also proposes a nightly export of ten years of distribution history to a separate data lake to feed the AI workloads.
Governance and Budget

Visit us at: https://www.examsempire.com/c-baipa-2610

The responsible-AI statement the board published to members last spring commits that no change to a member's payout is made without a named reviewer's approval. The budget committee will not approve a production commitment based on assumptions; it has asked for an estimate built from consumption measured on the configuration that will actually run in production. Royalty distribution periods and finance posting periods are defined differently today, and finance reconciles the two by hand each quarter; Peltonen wants any AI-facing royalty reporting to reconcile with finance without that manual step. Lehtonen has asked how the team would notice if the agent's behaviour shifted after go-live, and wants a solution blueprint the steering group can sign off within six weeks.

Your Assignment

You are the solution architect engaged by Kantele Rights Collective for a six-week design-validation phase on SAP Business AI Platform. The IT director, Aino Lehtonen, expects a solution blueprint covering a Joule-based member assistant and an AI agent for dispute triage. You will meet the head of member services, the finance director and finance core lead, the data protection lead, the budget committee, and the delivery contractor. Each stakeholder holds a position that is reasonable on its own terms. Your recommendations must respect the board's clean-core principle and its published responsible-AI commitment, the finance reconciliation requirement, and the budget committee's demand for a measured cost basis

1. CHALLENGE 1 — Deciding Where Dispute Tracking Lives Relative to the Finance Core

• Respond to the finance core lead's proposal to change standard ERP objects for dispute tracking, and recommend where the dispute-tracking capability should live.

• Advise the extension team on their prototype, which reads open-item data through an undocumented custom view a former implementation partner left in the ERP core.

• Explain your recommendation to the finance director in terms of update readiness and finance visibility, not tooling.

Checkpoints

• The recommendation addresses the finance director's visibility need and the board's clean-core principle together.

• The guidance on the prototype's data access is consistent with the board's clean-core principle.

2. CHALLENGE 2 — Restoring Confidence in the Agent's Proposals for Two Regional Broadcasters

• Diagnose why the agent's rights-holder proposals are wrong for two regional broadcasters before agreeing to any change to the agent itself.

• Respond to the contractor's offer of three point-to-point connectors in light of the IT director's preference for a repeatable integration approach and next year's new streaming sources.

• Define how the team should confirm that the corrective change resolves the pilot problem before the agent is reconnected.

Checkpoints

• The diagnosis is supported by the pattern of evidence from the pilot.

• The integration recommendation is defensible against the IT director's preference and next year's reporting changes.

• The proposed validation evidence is defensible to the IT director.

3. CHALLENGE 3 — Scoping the Member Assistant's Answers and Its Royalty Data Source

• Recommend what the Joule-based assistant may tell a member about works with several co-writers, reconciling the head of member services' request with the data protection lead's position.

• Respond to the contractor's proposal for a nightly export of ten years of distribution history to a separate data lake, considering the finance director's reconciliation requirement.

Checkpoints

• The assistant design can be explained convincingly to both the head of member services and the data protection lead.

Visit us at: https://www.examsempire.com/c-baipa-2610

• The data architecture recommendation answers the finance director's reconciliation concern.

4. CHALLENGE 4 — Moving the Dispute-Triage Agent From Pilot to Production

• Present the steering group with the order of activities that moves the dispute-triage agent from pilot to a limited production release.
• Respond to the head of member services' proposal to let the agent post royalty reassignments automatically to clear the backlog before the next distribution run.
• Explain to the budget committee how you will produce a cost estimate it can rely on, given that the pilot ran in an environment shared with developer experiments.
• Answer the IT director's question about how the team will notice a change in the agent's behaviour after go-live.

Checkpoints

• The rollout order is justified dependency by dependency.
• The response to the backlog proposal is consistent with what the board has published to members.
• The cost basis is one the budget committee can trace to its stated requirement.
• The answer on post-go-live behaviour explains how a change would become visible, and to whom.

Question: 1

CHALLENGE 1 — Deciding Where Dispute Tracking Lives Relative to the Finance Core

During the architecture workshop, the finance core lead proposes changing standard ERP objects: dispute fields added to the standard open-item records and an approval step coded into the standard payment run. Finance users already work there, and the change could be live within the current quarter. Dispute records today sit in the distribution engine and member-services spreadsheets, and the agent will soon generate more of them.

Which recommendation best fits Kantele's architecture principles while still meeting the finance director's visibility requirement?

A. Use in-app extensibility for a single dispute-status field that finance maintains by hand, and defer any dispute application until volumes justify one.
B. Build dispute tracking as a decoupled platform application that reads finance data through released interfaces, leaving standard objects untouched.
C. Build dispute tracking on the platform, but write dispute status back into the standard open-item records so existing finance reports display it.
D. Approve the standard-object changes for this quarter and add a backlog item to move them to the platform once dispute volume justifies an application.

Answer: B

Explanation:

Dispute data and workflow originate outside the ERP, in the distribution engine and soon the agent, and will change on their own lifecycle, which suits a decoupled extension on the platform. The platform application can show dispute status to finance users while reading finance data through released interfaces, so standard objects stay untouched, as the board's clean-core principle requires.

Question: 2

CHALLENGE 1 — Deciding Where Dispute Tracking Lives Relative to the Finance Core
The extension team's prototype on the platform reads open-item data through an undocumented custom view that a former implementation partner left in the ERP core as a workaround. They used it because the released interface did not appear to expose one remittance attribute they wanted.
What should the solution architect direct the team to do first?

A. Rebuild the custom view inside the core on top of released objects, so the dependency is cleaner while the extension's access path stays the same.
B. Continue on the custom view while asking SAP to expose the attribute in a future released interface, then switch over once it is delivered.
C. Replace the view with a direct read of the underlying standard tables from the extension, so no custom object needs to remain in the core.
D. Check for a released interface or an equivalent; if none exists, keep the view under a time-boxed exception the architecture board approves.

Answer: D

Explanation:
The first step is to confirm whether released interfaces can satisfy the requirement, because a custom core view is exactly the dependency the clean-core principle avoids. If a genuine gap remains, a time-boxed exception approved by the architecture board keeps the prototype moving while making the dependency visible and scheduled for removal, rather than silently inherited.

Question: 3

CHALLENGE 2 — Restoring Confidence in the Agent's Proposals for Two Regional Broadcasters
In the agent pilot, the dispute-triage agent proposes the wrong rights holder in roughly one case in six for two regional broadcasters. The contractor wants to retune the agent's matching instructions. Other broadcasters reporting the same regional repertoire through file transfers show normal accuracy.
What is the most likely root cause the architect should investigate before any change to the agent?

A. The intake route those two broadcasters use loses or reorders work identifiers, so the agent reasons over incomplete usage records, not the full report data.
B. The agent's instructions are too general for broadcast works, so it needs a separate instruction variant per broadcaster that encodes each report layout explicitly.
C. The rights-holder registrations in the distribution engine are outdated for those regions, so a bulk cleansing project must precede any further agent testing.
D. Those broadcasters' repertoire includes many co-written works, which the agent matches less reliably than works by a single composer or lyricist.

Answer: A

Explanation:
The two affected broadcasters are the ones whose emailed spreadsheets pass through an unmaintained reformatting script, while the same repertoire reported by file transfer matches normally, which points

to an upstream intake defect rather than an agent defect. An agent can only reason over the records it receives, so incomplete or reordered identifiers produce wrong proposals however well it is instructed.

Question: 4

CHALLENGE 3 — Scoping the Member Assistant's Answers and Its Royalty Data Source
The head of member services wants the Joule-based assistant to answer any payout question, including how earnings on co-written works were split. The data protection lead points out that each co-writer's individual payout amount is that member's personal financial data.
Which assistant design best satisfies both stakeholders?

A. Show co-writers' total quarterly payouts in rounded bands rather than exact figures, so members see each co-writer's scale without precise amounts.
B. Exclude all co-writer information from the assistant, so members asking about co-written works are always redirected to a member services agent by email.
C. Show co-writer figures to members who have opted in to detailed statements in the portal, applying that opt-in to every question they ask.
D. Limit the assistant to the asking member's own entitlements and each work's registered split percentages, excluding other members' overall payouts.

Answer: D

Explanation:
Members can see how a work's earnings were divided through the registered split percentages and their own entitlement, which answers most payout questions. Co-writers have already agreed those splits among themselves, so per-work shares are known to them; other members' overall payouts stay out of scope, which respects the data protection lead's position without gutting the assistant's usefulness.

Question: 5

CHALLENGE 4 — Moving the Dispute-Triage Agent From Pilot to Production
The steering group asks for the order in which the dispute-triage agent should move from pilot to production. Four candidate sequences are on the table.
Which sequence should the architect present?

A. Define which agent actions need human approval; estimate consumption from the pilot and agree the commercial model; validate outcomes on the managed feeds; make a limited production release.
B. Define which agent actions need human approval; make a limited production release; validate outcomes on the managed feeds; estimate consumption and agree the commercial model.
C. Define which agent actions need human approval; validate outcomes on the managed feeds; estimate consumption and agree the commercial model; make a limited production release.
D. Define which agent actions need human approval; validate outcomes on the managed feeds; make a limited production release; estimate consumption and agree the commercial model.

Answer: C

Explanation:
Approval rules must exist before validation so the validation tests the agent as it will actually operate, with payout-affecting actions routed to reviewers. Validation then produces the measured consumption the budget committee requires for the estimate, and only an approved budget allows the limited release.
Topic: 2

Question: 6

Aceros del Bajío, a Mexican structural-steel fabricator running SAP S/4HANA Cloud, receives roughly 400 incoming steel coils a week. Inspectors at its three plants will record coil test results on tablets during each shift. A fit-to-standard workshop confirmed that the standard inspection apps cannot capture the plant-specific test sequence, and the architecture board has approved a side-by-side extension on SAP Business AI Platform that consumes only released APIs, in line with the company's clean-core commitment.
Three requirements remain for the build plan. The inspection rule set is owned by the quality team, changes monthly, and must be updated by that team without waiting for a developer release. When a test fails, the ERP system must receive a quality notification and the affected batch must be blocked from production use. Results are posted throughout each shift, and coil stock is drawn for production continuously during every shift. Each plant has one quality manager, who works day shifts only. The ERP system exposes released APIs for creating quality notifications and for changing batch status.
The lead developer has drafted four three-part build plans and asks you, as the solution architect, to approve the one that meets every requirement.
Which build plan should you approve?

A. Hold the rule set as configurable data in the extension, maintained by quality; collect failed results during the shift; at shift end, call the released APIs to create notifications and block batches.
B. Keep the rule set in a quality-owned spreadsheet loaded into the extension at each monthly deployment; on a failed test, call the released notification API; in the same processing step, block the batch via the released batch-status API.
C. Hold the rule set as configurable data in the extension, maintained by quality; on a failed test, call the released notification API; in the same processing step, block the batch via the released batch-status API.
D. Hold the rule set as configurable data in the extension, maintained by quality; on a failed test, call the released notification API; block the batch via the released batch-status API once the plant quality manager approves.

Answer: C

Explanation:
This plan meets every requirement with nothing added. The rule set lives as configurable data the quality team can change monthly without a developer release, and both outcomes required on a failed test, the quality notification and the batch block, are executed through the ERP system's released APIs at the moment the result is recorded. The design stays within the approved clean-core approach, and the block takes effect before production can draw the affected stock.

Question: 7

Straits Meridian Bank in Singapore has piloted an AI agent, built with Joule Studio and monitored through SAP AI Agent Hub, that reviews SME customers' cash-flow data and proposes changes to their overdraft limits. During the pilot, proposals reached relationship managers as emailed spreadsheets, and approved changes were re-keyed manually into the lending system, which accounted for most of the two-day turnaround. For production, the business sponsor wants the agent to apply limit changes directly, cutting turnaround from two days to minutes, and the agent's technical user has already been granted the write authorization to do so.
The model-risk team has two findings. First, the bank's AI policy, aligned with its regulator's expectations, requires a named human to remain accountable for any decision that changes a customer's credit exposure. Second, a sample review found the agent's proposals well reasoned in 94% of cases; in the rest it misread seasonal revenue dips as distress and proposed reductions.
The business sponsor argues that 94% accuracy plus full logging of every agent action is sufficient oversight. The CIO wants a production design the bank can extend to other credit products without re-architecting, and the head of SME banking wants turnaround to improve noticeably for customers.
Which production design should you recommend?

A. Keep the agent's technical user read-only and post each proposal to the manager's workflow inbox for keying into the lending system, because the agent then never writes to credit exposure itself.
B. Let the agent apply changes directly, name the account's relationship manager as accountable owner, and send a same-day digest with one-step reversal, because each change keeps a named owner.
C. Let the agent apply limit increases directly and route only proposed reductions to a relationship manager, because increases carry little customer-harm risk and the sampled errors in the pilot all involved reductions.
D. Keep the agent proposing changes and route each to the account's relationship manager for approval in the workflow before the agent executes it, because the agent still prepares every change.

Answer: D

Explanation:
An approval step inside the workflow satisfies the policy as written: every change to credit exposure has a named, accountable relationship manager, while the agent still does the analysis and executes the approved change. Approving inside the workflow removes the spreadsheet hand-off and the manual re-keying that consumed most of the two days, so turnaround drops to the time a manager needs to review a prepared proposal. The same propose-approve-execute pattern can be reused for the bank's other agents and credit products, and the human review catches the seasonal misreads before they reach a customer.

Question: 8

You are the solution architect advising Nordhavn Freight, a Danish logistics group running ports and road haulage across the Nordics. Its integration estate runs on an on-premise middleware platform whose vendor support ends in 16 months; the vendor will extend support by six months at most, at a premium. The CIO has mandated SAP Integration Suite as the single target integration capability for all new and migrated flows.
An inventory lists 142 interfaces. The integration lead notes that many were built project by project: at least 35 move the same shipment-status data between the same systems through separate point-to-

point flows, and a separate group of about 20 serve applications scheduled for decommissioning within the year. The team has four integration developers, and the CFO has capped the migration budget at a level the integration lead estimates covers about 90 interface rebuilds.
Three proposals are on the table. The integration lead favours a one-for-one rebuild of every flow so nothing is missed. The infrastructure manager suggests negotiating extended middleware support and moving flows only as projects touch them. An enterprise architect wants to apply the Integration Solution Advisory Methodology before committing to any migration plan.
Which recommendation best meets the support deadline, the budget cap, and the CIO's goal of a governed integration estate?

A. Rebuild all 142 interfaces one-for-one on SAP Integration Suite in prioritised waves, because a like-for-like migration minimises functional risk and finishes predictably inside the 16-month support window.
B. Negotiate extended middleware support and move each interface only when a project next changes it, because this spreads cost across project budgets and avoids a programme the team cannot staff.
C. Classify the interfaces by pattern with the methodology, retire and merge redundant flows, and migrate the rest in waves before support ends, because the migration then covers only flows still needed.
D. Use the methodology to target the 35 duplicated shipment-status flows, merge them into one reusable flow, and rebuild every other flow one-for-one in waves, because duplication is the main gap.

Answer: C

Explanation:
This applies the advisory methodology where it creates value: assessing the inventory against the application roadmap and classifying flows by pattern shows which flows can be retired and which duplicates can collapse into reusable flows. Retiring the roughly 20 decommissioned-application flows and merging the 35 duplicates brings the rebuild scope down to roughly the 90 the budget covers, so the four-person team can migrate in waves before support ends, and the result is a governed, pattern-based estate on SAP Integration Suite rather than a copy of the old one.

Question: 9

Kestrel Ridge Water, a Canadian municipal water utility, is finalising the commercial plan for its SAP Business AI Platform programme. The programme has two parts. The first is a set of integration flows and extension applications that will run steadily from go-live, with volumes the utility can forecast from five years of meter and work-order history. The second is an AI agent pilot for field-crew scheduling whose usage could range from a dozen supervisors to all 600 field staff within the first year, depending on pilot results.
The account team has offered two commercial options that can be used separately or together. Under the fixed subscription, named services are bought at agreed quantities for the contract year; this offer allows no mid-year quantity changes, which are made only at renewal. Under the usage-based option, services draw down a pre-committed credit balance as they are consumed, and credits can be topped up at any point in the contract year at the committed rate. The quote prices the steady workloads lower under fixed subscription than under usage-based credits at forecast volumes. After the steady workloads, the approved envelope leaves room for pilot usage equivalent to about 150 users.

The CFO requires annual platform spend to stay inside that envelope. The CIO does not want commercial terms to hold the pilot back if adoption grows mid-year. The finance team prefers a single model to keep administration simple.
Which commercial recommendation best balances the CFO's spending envelope with the CIO's need to scale the agent pilot?

A. Subscribe the steady workloads at fixed quantities and run the pilot on usage-based credits tracked monthly against the envelope, because each workload is bought to match its predictability.
B. Place the whole programme on the usage-based credit model, because one flexible pool lets the agent pilot scale freely and keeps administration as simple as the finance team has asked.
C. Subscribe the steady workloads now and add the pilot to the subscription at a dozen-supervisor quantity raised at renewal, because the pilot starts small and no unused capacity is paid for.
D. Place the whole programme on fixed subscription with the pilot sized at a mid-range estimate of 300 users, because the CFO gets one predictable annual figure and the pilot has headroom to grow.

Answer: A

Explanation:
Each workload is placed on the model that suits its predictability. The forecastable workloads take the lower fixed-subscription price shown in the quote, while the pilot draws on the usage-based credit balance, so its consumption can rise with adoption through mid-year top-ups at the committed rate. Monthly tracking of credit consumption gives the CFO a monthly control point, so pilot growth is funded within the envelope and any case for going beyond it is decided on usage evidence, which balances the two constraints better than any single-model plan.

Question: 10

Seguros Caravela, a Brazilian property and casualty insurer, is piloting an AI agent built with Joule Studio to support its claims adjusters. The agent reads a claim, checks it against the policyholder's contract terms, and drafts a settlement recommendation that an adjuster reviews before anything is paid.
After three weeks the claims director reports mixed results. The drafts are fluent and correctly structured, and the agent retrieves claim headers and loss details from the claims system without errors. However, in 11 of 40 sampled cases the recommendation applied general market practice on deductibles and exclusions instead of the clauses in Seguros Caravela's own policy wordings, which differ by product line and issue year. Adjusters had to rewrite those drafts.
A configuration review confirms that the agent's instructions tell it to verify coverage against the applicable policy and that the claims-system connection is healthy. The agent's configured tools are a claims-system lookup and a recommendation-document template. Policy wordings are maintained by the underwriting department as versioned documents per product line and issue year, and underwriting issues revised wordings for each new issue year. The claims director wants the defect fixed before the pilot expands to two more regions, without exposing more personal data than the pilot already uses.
As the solution architect, which corrective action addresses the root cause of the defective recommendations?

A. Give the agent governed access to underwriting's versioned policy wordings as reference content, selected by each claim's product line and issue year, so its coverage check uses the applicable clauses.

B. Add a mandatory step in which the adjuster confirms the applicable policy clauses before the agent drafts, so every recommendation is anchored to the correct wording for that claim's product line.
C. Commission a custom-trained model on five years of settled claims and adjuster notes, so the agent learns how Seguros Caravela's adjusters have interpreted coverage across all of its product lines.
D. Expand the agent's instructions with a detailed summary of deductible and exclusion rules per product line and issue year, so the agent applies Seguros Caravela's own terms consistently whenever it drafts a recommendation.

Answer: A

Explanation:
The evidence isolates the gap: the claims connection works and the instructions ask for a coverage check, but the agent's tools give it only claim data and a document template, with no access to the policy wordings, so it falls back on general knowledge whenever a clause is needed. Grounding the agent in the insurer's own versioned policy wordings, selected by product line and issue year, gives the coverage check the contract text it must reason over. This fixes the root cause at the context layer and uses contract documents rather than additional personal data.

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/