RapidCanvas

PARTNER CASE STUDY | RAPIDCANVAS + GOOGLE CLOUD

Gemini and Google Cloud, From Pilot to Production

How RapidCanvas puts Gemini, Cloud SQL, BigQuery, Cloud Storage and Google Workspace to work inside live customer deployments, from document extraction to production databases.

12
use cases

7
in production

September 2026
edition



Real AI transformation, accelerated.

AT A GLANCE

RapidCanvas + Google Cloud

RapidCanvas is a Managed AI Execution Company: production systems, not slide decks. On Google Cloud that is a dozen live, in-delivery, piloted or proof-of-concept deployments, each reading governed data in BigQuery, Cloud SQL, Cloud Storage, Drive or Sheets, and calling Gemini for the extraction behind the result. Our Hybrid Approach™, an agentic platform paired with expert-in-the-loop review, is why it ships fast and keeps running.

Canvas pipelines move data through validation gates. Data Apps give reviewers a screen instead of a spreadsheet. AI Agents and Ask AI turn Gemini calls into assistants and decisions, routed through the model gateway with rate limiting and cost accounting. Google Cloud teams should read this as consumption RapidCanvas drives directly: more Gemini calls, more Cloud SQL and BigQuery reads, more Cloud Storage traffic.

12
Google Cloud use cases documented

7
use cases already in production

10
Google Cloud and Workspace services in use

PARTNERSHIP
RapidCanvas is a Google Cloud Partner listed in the Partner Directory, engaging customers directly today; a Google Cloud Marketplace listing is on the roadmap.

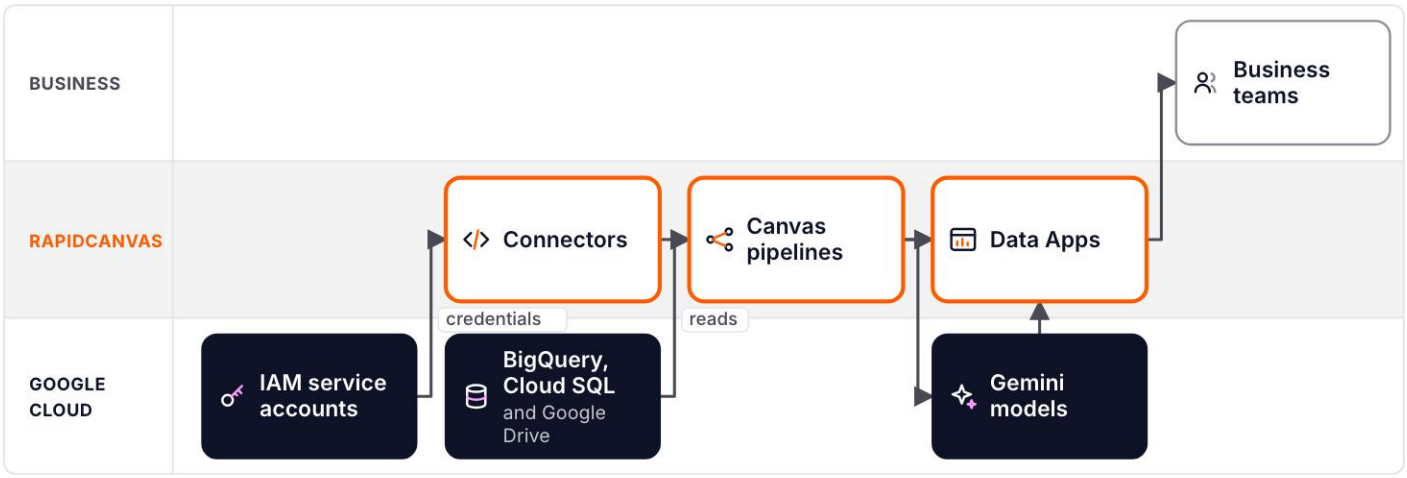
When RapidCanvas Is a Fit

1. A customer has governed data in BigQuery, Cloud SQL, Drive or Sheets and wants agents built.
2. Gemini usage is stalling on quota, rate limits or inconsistent JSON output at volume.
3. A workload needs a managed Cloud SQL or BigQuery database behind a hosted application.
4. A team wants Gemini extraction with human review, not an unsupervised model deciding alone.

REFERENCE ARCHITECTURE

How RapidCanvas Builds on Google Cloud

The same shape recurs across every use case: scoped service-account credentials, data from Google Cloud, Gemini behind a validation layer, and RapidCanvas assembling the result for a business user.



Credentials stay scoped

IAM service accounts, scoped to the project, held in RapidCanvas Secrets rather than in code or config files.

Data stays in place

BigQuery, Cloud SQL, Cloud Storage and Drive are read where the customer already governs them, not copied elsewhere first.

Models sit behind validation

Gemini output is schema-checked, confidence-scored and, below threshold, routed to a person before anything is written.

Google Cloud Services in These Stories

Gemini
Structured extraction, classification and reasoning, behind a validation layer

Cloud SQL
Application and warehouse database for RapidCanvas-hosted services

BigQuery
Analytics source and change-data-capture destination for usage reporting

Cloud Storage
Signed-URL document access, video uploads, model artifacts and backups

Google Drive
Document intake and automated filing through a service account

Google Sheets
Operational trackers and benchmarks business teams already keep

Cloud Run
Public webhook proxy fronting a RapidCanvas-hosted assistant

Compute Engine
GPU virtual machines running vision-model inference and training

Datastream
Change-data-capture replication from a production database into BigQuery

Google Maps Platform
Location discovery for go-to-market signal scoring

IN THIS DOCUMENT

12 Use Cases, 10 Google Cloud Services

Each use case runs on native Google Cloud services, with RapidCanvas building the pipelines, agents and Data Apps on top.

NO.	USE CASE	INDUSTRY	STATUS
01	Custody Statement Reconciliation Runs on Gemini	Private wealth management, Latin America	**In production**
02	Five Payroll Documents, Parsed and Joined in Cloud SQL	Construction compliance, North America	**In production**
03	Grounded Gemini Signals Prioritize New Sales Accounts	Electronics manufacturing, global	**In production**
04	Roadside Insurance Cards Read by Gemini When Codes Fail	Public infrastructure, Africa	**In production**
05	Dashcam Video Meets GPU Inference on Google Cloud	Public infrastructure, Africa	**In production**
06	Chat Analytics Flow From BigQuery Into Cloud SQL	Professional services, global	**In production**
07	A Five Schema Lakehouse Built on Cloud SQL	Real estate development, Latin America	**In production**
08	Wealth Documents Classified and Extracted by Gemini	Wealth management technology, North America	**Pilot**
09	Restaurant Invoices Move From Drive Into a Ledger	Hospitality and restaurants, global	**In delivery**
10	A Test Coverage Assistant Reads a Live Sheet	Enterprise software, global	**In delivery**
11	A Gemini Pipeline Scopes Bids for Estimators	Construction estimating, North America	**Proof of concept**
12	Structured Extraction and Managed Databases, Repeated	Multiple industries, global	**Reference pattern**

USE CASE 01 | PRIVATE WEALTH MANAGEMENT, LATIN AMERICA ● IN PRODUCTION

Custody Statement Reconciliation Runs on Gemini

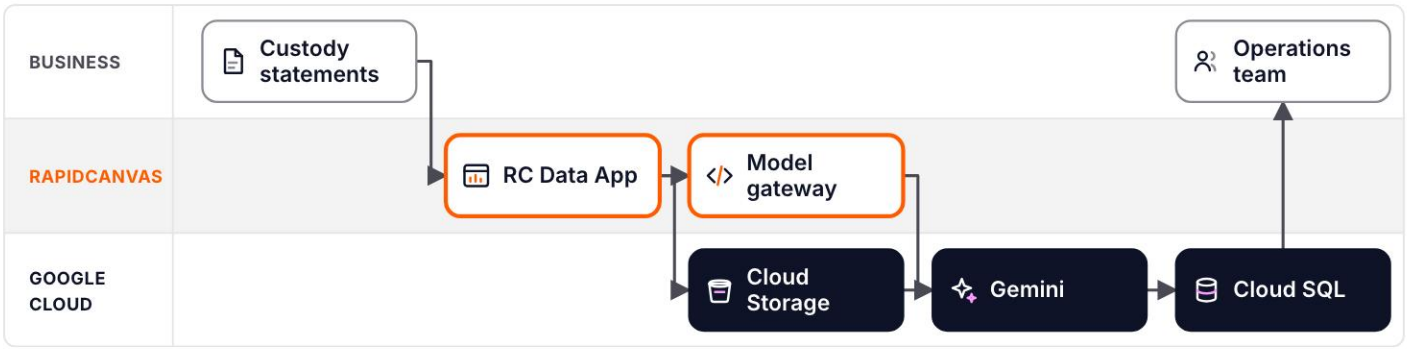
Gemini | Cloud SQL | Cloud Storage | Canvas pipelines | Data Apps | Model gateway | Expert-in-the-Loop

The Challenge

A multi-family office receives custody and bank statements from dozens of institutions, in a different layout each time, for onshore and offshore portfolios. Positions must be matched to the office's portfolio system to catch missing holdings, wrong quantities and stale prices before client reporting. Manual reconciliation did not scale, and the starting model tier ran out of quota within weeks.

Solution Implemented

Users upload statements through a Data App, where fuzzy matching ties each one to the right portfolio. Extraction never uploads the file to the model: the backend gets a signed Cloud Storage URL and calls Gemini through the model gateway with a fixed seed. Deterministic rules match identifiers first, with a Gemini fallback for the rest. Views classify every row, prices follow domain rules, and results post to the customer's systems, recorded in Cloud SQL.



How It Works

1. Upload statements in the Data App; fuzzy matching ties each one to an account or pauses.
2. Get a signed Cloud Storage URL and call Gemini via the gateway with a fixed seed.
3. Run deterministic rules first, then a Gemini fallback under a strict JSON schema for the rest.
4. Classify rows, apply price rules, and write results to Cloud SQL and the customer's own system.

Results and Benefits

- Live since 2026, with document onboarding in ramp.
- A rate limiter cleared an incident that had stalled a small share of a large queue.
- Model lineage moved through three Gemini tiers without changing the pipeline.

HOW IT RUNS ON GOOGLE CLOUD

Every extraction and match call runs on Gemini; Cloud Storage and Cloud SQL hold every document and result. Consumption grows with statement volume, and data stays in the customer's own systems.

USE CASE 02 | CONSTRUCTION COMPLIANCE, NORTH AMERICA ● IN PRODUCTION

Five Payroll Documents, Parsed and Joined in Cloud SQL

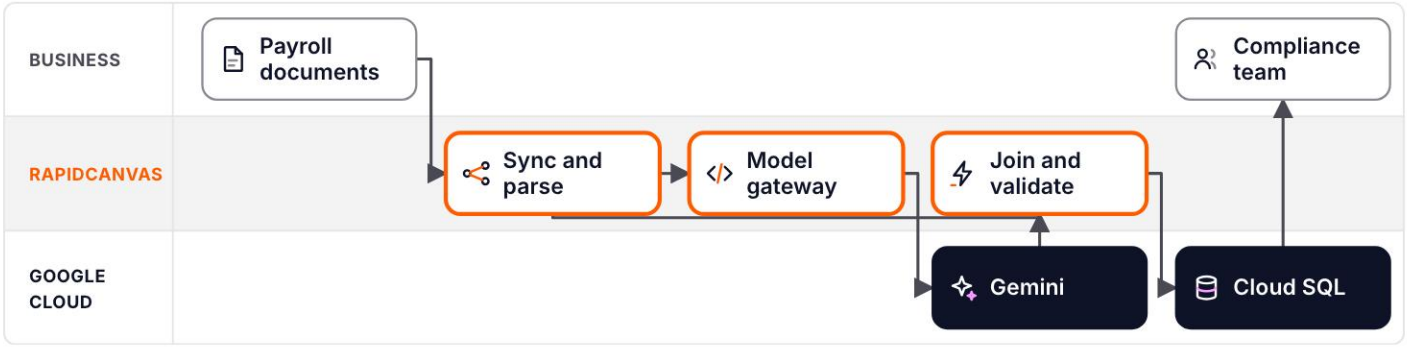
Gemini | Cloud SQL | Canvas pipelines | Data Apps | Model gateway

The Challenge

On public-works projects a general contractor is accountable for its subcontractors' labour compliance. Certified-payroll hours, timesheet hours and paystub pay must agree, apprentices must be paid the right share of the journeyworker rate, and daily apprentice ratios must hold. Documents arrive as PDFs in dozens of layouts for many subcontractors at once, which made manual checking slow.

Solution Implemented

A sync recipe mirrors each project folder into an artifact, and Gemini labels supporting documents so folder rules never need regex. Parse recipes send certified payrolls, timesheets and paystubs as raw PDFs, in parallel chunks, to Gemini through the model gateway. A join step matches documents by name, with Gemini reconciling names that differ, and computes expected pay from hours, rate and apprentice share. Validation flags discrepancies, and results replace each project's Cloud SQL rows for a Data App.



How It Works

1. Mirror project folders into an artifact; Gemini labels documents so rules never need regex.
2. Parse certified payrolls, timesheets and paystubs as raw PDFs through Gemini via the gateway.
3. Join documents by name, with Gemini reconciling names, and compute expected pay.
4. Flag ratio and pay discrepancies, write results to Cloud SQL, shown in a payroll Data App.

Results and Benefits

- Live for the compliance team since 2026.
- A full project run takes roughly thirty minutes to two hours depending on volume.
- Compliance staff review flagged discrepancies in a dashboard instead of checking by hand.

HOW IT RUNS ON GOOGLE CLOUD

Gemini does five-document extraction and name reconciliation, and Cloud SQL holds every payment and ratio table. The Canvas pipeline runs inside Google Cloud, reaching the database directly.

USE CASE 03 | ELECTRONICS MANUFACTURING, GLOBAL ● IN PRODUCTION

Grounded Gemini Signals Prioritize New Sales Accounts

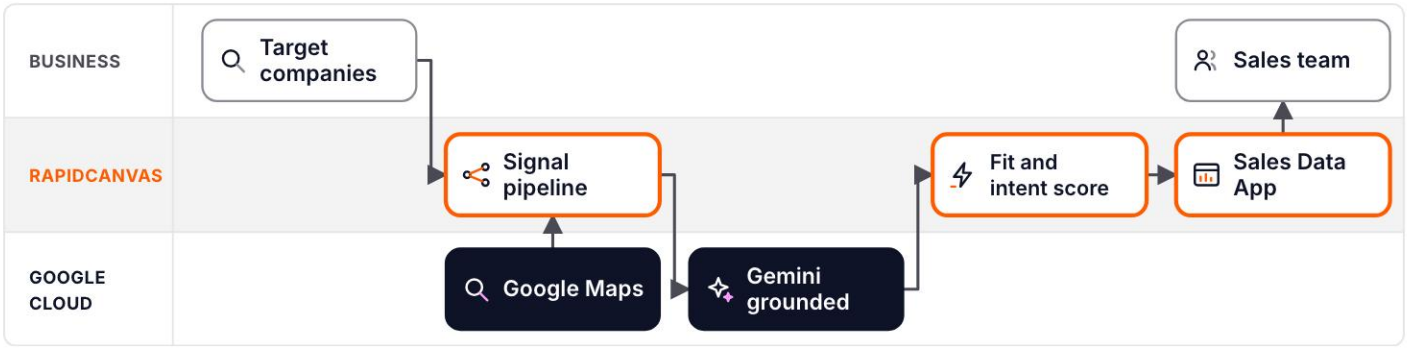
Gemini | Google Maps Platform | Canvas pipelines | Data Apps | Model gateway

The Challenge

A global electronics manufacturer's sellers needed to know which companies in two fast-moving verticals are building, hiring or clearing regulatory milestones, with evidence they could cite, refreshed weekly. Generic web search diluted results across regions and surfaced stale news, and every signal had to map to a site location and the manufacturer's product portfolio before a seller could act.

Solution Implemented

Canvas pipelines build a location-level picture of each company from job boards and Google Maps, then a signals node runs region-split Gemini queries with Google Search grounding, covering manufacturing, research, regulatory approvals and capital spending over a two-year lookback. Grounding metadata is harvested so every signal carries citable URLs, and a freshness control drops stale items before scoring. Gemini also classifies whether a company operates infrastructure the manufacturer sells into. Fit and intent scores queue companies in a Data App.



How It Works

1. Build a location-level picture of each company from job boards, career pages and Google Maps.
2. Run region-split Gemini queries per company with Google Search grounding and citation harvesting.
3. Drop stale signals with a freshness control, then classify companies against portfolio fit.
4. Score fit and intent, and route leads into a sales Data App for the account team.

Results and Benefits

- Live since 2026 as a tier-one account.
- One run discovered hundreds of companies from an initial seed list, qualifying most of them.
- The grounded-signals node covers several companies in under an hour.

HOW IT RUNS ON GOOGLE CLOUD

Every signal and narrative runs on Gemini with Google Search grounding, and location discovery runs on Google Maps Platform, scaling with the number of target companies a seller adds.

USE CASE 04 | PUBLIC INFRASTRUCTURE, AFRICA ● IN PRODUCTION

Roadside Insurance Cards Read by Gemini When Codes Fail

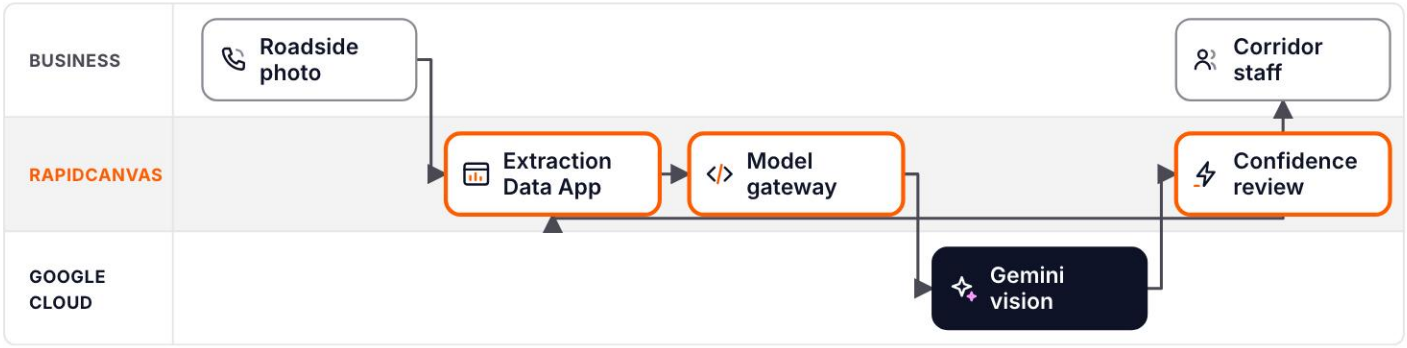
Gemini | Data Apps | Model gateway | Expert-in-the-Loop

The Challenge

Trucks crossing a regional trade corridor carry an insurance certificate whose QR code is the machine-readable key to enrolment. In the field the code is often unreadable, leaving staff to transcribe faint identifiers from glared, stamped phone photos by hand. The chassis number and registration must be right, and the authority needed a reliable fallback when the code fails.

Solution Implemented

The enrolment system calls a hosted endpoint with one image when a QR read fails, or staff drop a batch into a Data App. A library renders the image, applies a bilingual prompt with character-ambiguity rules and a fourteen-field schema, and calls Gemini through the model gateway; a second pass re-reads the chassis and registration fields. Fields are reshaped to value and confidence, flagged for review below 0.85, and appended to an artifact the Data App lists.



How It Works

1. Receive one image, or a batch through the Data App, when a roadside QR read fails.
2. Render the image and call Gemini with a layout-aware prompt and a fourteen-field JSON schema.
3. Re-read the chassis and registration fields in a second pass with a confidence per field.
4. Flag any identity-critical field below 0.85 for review, and append the record to an artifact.

Results and Benefits

- Deployed on the authority's production tenant.
- The chosen model was benchmarked at one hundred percent on the internal sample set.
- Every extraction returns its own cost, so the authority sees what a card costs to read.

HOW IT RUNS ON GOOGLE CLOUD

Every extraction call runs on Gemini through the model gateway, with cost tracked per card. The pattern gives Google Cloud a field-operations use case for Gemini vision.

USE CASE 05 | PUBLIC INFRASTRUCTURE, AFRICA ● IN PRODUCTION

Dashcam Video Meets GPU Inference on Google Cloud

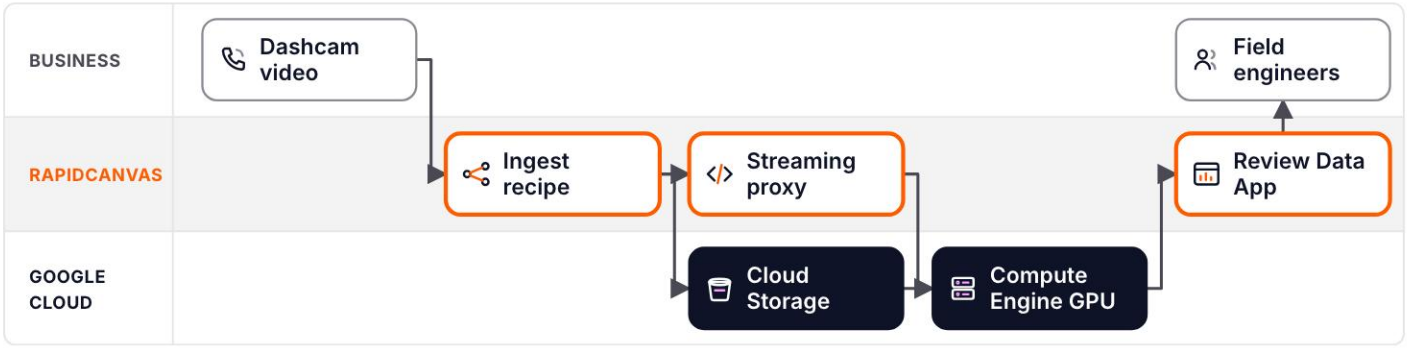
Compute Engine | Cloud Storage | Canvas pipelines | Data Apps | Connectors

The Challenge

A road authority assesses corridor defects by driving and inspecting segments manually. It wanted defects identified from vehicle video with GPS-tagged evidence and repair-budget estimates, while cutting field visits. Inference on a long video takes tens of minutes on a GPU, longer than any platform request should hold open, so the detector had to run outside the main application.

Solution Implemented

Field videos are pulled by recipes, and frames and GPS metadata are extracted. A GPU machine on Compute Engine, outside the platform, runs a vision detector as a service; a hosted proxy streams each video and relays results back, uploading large videos to Cloud Storage first. Model weights come from a registry under a champion alias, so promoting a model is a configuration change. Reviewers reclassify detections in a Data App, and a recipe uploads labelled frames.



How It Works

1. Pull field videos and extract frames and GPS metadata with RapidCanvas recipes.
2. Stream video to a Compute Engine GPU virtual machine and relay detections back.
3. Let reviewers confirm or reclassify detections in a Data App, synced to an artifact.
4. Upload labelled frames to Cloud Storage for retraining, and promote a model by alias.

Results and Benefits

- Live since 2026, with a defect-class enhancement soon after.
- Success criteria set pipeline uptime and a three to six hour video-to-dashboard latency target.
- The authority reports significantly fewer field visits.

HOW IT RUNS ON GOOGLE CLOUD

GPU inference and training run on a Compute Engine virtual machine, and Cloud Storage holds every large video and frame. Long-running inference stays on infrastructure the authority already pays for.

USE CASE 06 | PROFESSIONAL SERVICES, GLOBAL ● IN PRODUCTION

Chat Analytics Flow From BigQuery Into Cloud SQL

BigQuery | Cloud SQL | Datastream | Canvas pipelines | Data Apps | Connectors

The Challenge

Administrators of an internal knowledge assistant needed to see who uses it, how well it answers, and what fails, without hammering the live chat service. A first design that queried the chat database directly saturated the shared instance once a handful of people read it at once, so analytics had to become cheaper and independent of the chat store.

Solution Implemented

Datastream replicates the chat service's database into BigQuery in near real time, exposed as a view. Every thirty minutes an analytics API reads that view with the BigQuery client, maps its schema to indexable Cloud SQL columns, and streams a parameterised query in pages. Rows are staged into Cloud SQL, merged on chat id, with deletions soft-applied; a legacy source stays one switch away as a rollback path. A materialised fact view unpacks feedback for an admin Data App.



How It Works

1. Replicate the chat database into BigQuery in near real time with Datastream.
2. Read the BigQuery view every thirty minutes and map its schema to Cloud SQL columns.
3. Stage and merge rows into Cloud SQL on chat id, keeping a legacy source as rollback.
4. Materialise a fact view and serve usage and quality endpoints in an admin Data App.

Results and Benefits

- Live since 2026.
- A full sync completes in about ten seconds every half hour, with lag under a minute.
- A shared-instance CPU incident was closed by the materialised-facts design.

HOW IT RUNS ON GOOGLE CLOUD

Datastream and BigQuery replace direct reads against a production database, and Cloud SQL hosts the analytics store, showing how change-data capture protects a live service.

USE CASE 07 | REAL ESTATE DEVELOPMENT, LATIN AMERICA ● IN PRODUCTION

A Five Schema Lakehouse Built on Cloud SQL

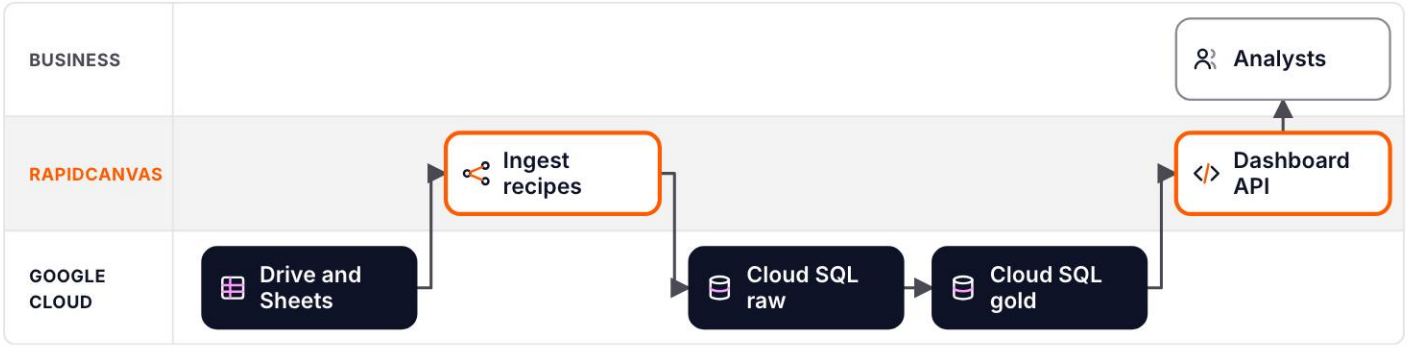
Cloud SQL | Google Drive | Google Sheets | Canvas pipelines | Data Apps | Connectors

The Challenge

A residential real-estate developer's market-intelligence app depended on a legacy schema of sixty tables that mixed application state with pipeline output and had no lineage. Benchmark spreadsheets and forms lived in Drive and Sheets, vendor files elsewhere. The team wanted a governed lakehouse with fail-closed data quality on the managed database it already ran.

Solution Implemented

One Cloud SQL for MySQL instance hosts five schemas: raw ingestion manifests, silver tables per source, a gold star schema, an operations schema for data-quality results, and the application schema. Benchmark recipes load a service-account key, list a Drive folder, and download only changed spreadsheets by comparing modified time and checksum against a manifest table. Declarative contracts validate every source at bronze, blocking critical failures; gold dimensions rebuild through shadow tables and atomic swaps.



How It Works

1. Host five schemas, raw, silver, gold, ops and app, in one Cloud SQL instance.
2. Download only changed benchmark files from Drive and Sheets by comparing time and checksum.
3. Validate every source at bronze against contracts, blocking partitions on critical failures.
4. Rebuild gold dimensions through shadow tables and swaps, and serve them to a dashboard.

Results and Benefits

- Live since 2026.
- The application reads a documented gold layer of several dimensions and facts.
- More than a dozen ingest recipes run with data-quality gates, and reports go out by email.

HOW IT RUNS ON GOOGLE CLOUD

Cloud SQL hosts the entire medallion lakehouse, and Drive and Sheets feed it as governed benchmark sources. A managed relational database can carry a full lakehouse design.

USE CASE 08 | WEALTH MANAGEMENT TECHNOLOGY, NORTH AMERICA ● PILOT

Wealth Documents Classified and Extracted by Gemini

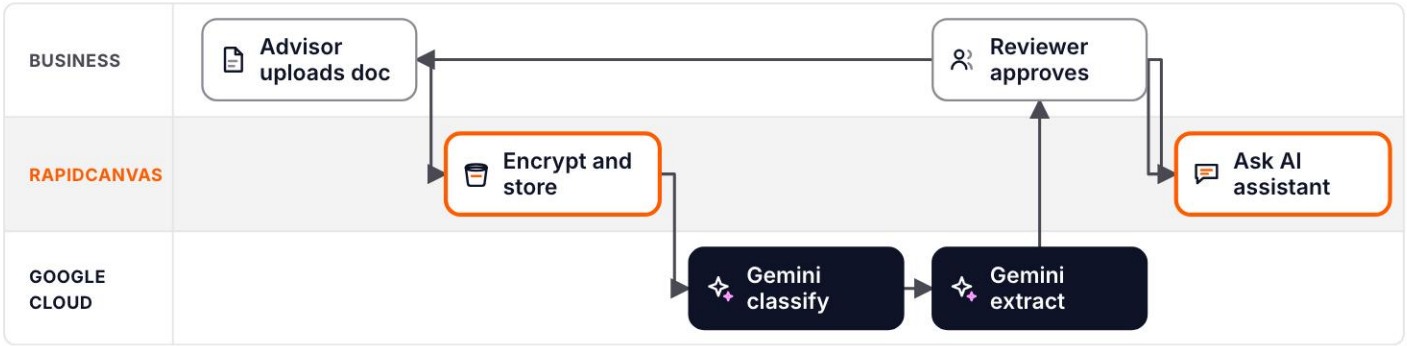
Gemini | Data Apps | Ask AI | Expert-in-the-Loop

The Challenge

A wealth-management technology firm's advisors answered client questions by hand-searching financial statements, tax forms, capital-call notices and property statements. An analysis of standard advisor questions found roughly forty percent answerable from documents on hand, if those were classified, entities extracted with ownership rules, and results reviewed before feeding a knowledge assistant.

Solution Implemented

An advisor uploads a PDF or spreadsheet, encrypted and stored in an artifact. A background task sends the first page to Gemini in JSON mode for one of eight document types with a confidence and reasoning; below 0.90 the full document goes through a second pass. A parser routes PDFs and spreadsheets through the same path. For extractable types, Gemini extracts entities with ownership rules, validated before reviewers approve, edit or reject them.



How It Works

1. Encrypt and store an uploaded document, then classify its type from one page with Gemini.
2. Escalate to a full-document second pass only when classification confidence falls below 0.90.
3. Extract entities with domain ownership rules, and validate output against a fixed enumeration.
4. Let reviewers approve, edit or reject entities, then feed approved ones into an assistant.

Results and Benefits

- Several delivery milestones shipped in 2026, with self-service upload in progress.
- An analysis found roughly forty percent of standard advisor questions answerable from documents on hand.

HOW IT RUNS ON GOOGLE CLOUD

Every classification and extraction call runs on Gemini, with model choice and thresholds set by configuration. Consumption grows with document volume as self-service upload rolls out.

USE CASE 09 | HOSPITALITY AND RESTAURANTS, GLOBAL ● IN DELIVERY

Restaurant Invoices Move From Drive Into a Ledger

Google Drive | Gemini | Cloud SQL | Canvas pipelines | Data Apps | Model gateway

The Challenge

A hospitality group's accountants received vendor invoices as PDFs in a per-restaurant folder, extracted line data by hand, coded each line to a ledger category, reconciled against the accounting ledger, and filed the PDF into a year, month and vendor folder before posting bills. The Drive tree had to stay as the team already used it.

Solution Implemented

A scheduled recipe walks each restaurant's intake folder with a service account, de-duplicates by content hash, and downloads new PDFs into an artifact. An extraction recipe sends each PDF to Gemini through the model gateway and writes structured invoice JSON. Categorisation runs in tiers: operator mappings, library matches, fuzzy retrieval, a closed-list Gemini pick, then human review, with the model only choosing from the restaurant's real category list. A load recipe upserts invoices into Cloud SQL.



How It Works

1. Walk each restaurant's Drive folder, de-duplicate by content hash and download new invoice PDFs.
2. Send each invoice to Gemini through the model gateway and write structured line-item JSON.
3. Categorise lines in tiers, from operator mappings to a closed-list Gemini pick with review.
4. Upsert invoices into Cloud SQL, let operators approve lines, and file the PDF back to Drive.

Results and Benefits

- In delivery since 2026, with user-acceptance and production deploys on push.
- A freshness alarm watches the Drive ingest, and every correction improves the next run.

HOW IT RUNS ON GOOGLE CLOUD

Gemini extracts every invoice line, Google Drive is the intake and filing destination, and Cloud SQL holds the ledger, keeping the customer's Drive structure as the system of record.

USE CASE 10 | ENTERPRISE SOFTWARE, GLOBAL ● IN DELIVERY

A Test Coverage Assistant Reads a Live Sheet

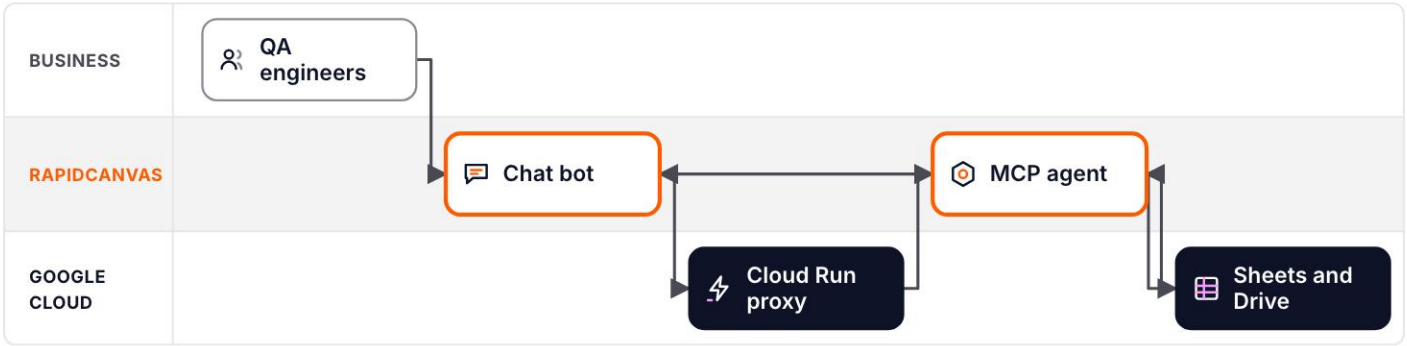
Google Sheets | Google Drive | Cloud Run | AI Agents | Connectors | Ask AI

The Challenge

An enterprise-software vendor's QA organisation kept asking the same questions: is this scenario covered, what failed last run, where are the blockers this week. Answers lived in a human-edited spreadsheet that changed shape without notice, a test corpus, and per-run files on a shared drive. The team wanted cited chat answers without giving a bot write access.

Solution Implemented

An MCP server reads the coverage sheet with a read-only scope, projecting known tabs into typed test cases or reading tabs generically from their header row after a rename. A Drive adapter reads run metadata from the shared drive by file property, contract-tested to expose no mutating method. Tools answer coverage questions with citations and an authority table. A chat bot runs a tool-use loop against the server, reached through a Cloud Run function.



How It Works

1. Read the coverage sheet with a read-only scope, tolerating tabs renamed or re-columned live.
2. Read run indexes and metadata from the shared drive by file property, with no write access.
3. Answer coverage and blocker questions with citations and an explicit source-authority table.
4. Reply in chat through a bot that runs a tool-use loop behind a Cloud Run proxy.

Results and Benefits

- In delivery, with the first use case live since 2026.
- More than a dozen tools serve several data sources.
- A live sheet change that renamed a tab and added columns did not break the assistant.

HOW IT RUNS ON GOOGLE CLOUD

Google Sheets and Drive stay the system of record, read only, while Cloud Run fronts the chat integration, feeding an AI agent without data leaving the customer's account.

USE CASE 11 | CONSTRUCTION ESTIMATING, NORTH AMERICA ● PROOF OF CONCEPT

A Gemini Pipeline Scopes Bids for Estimators

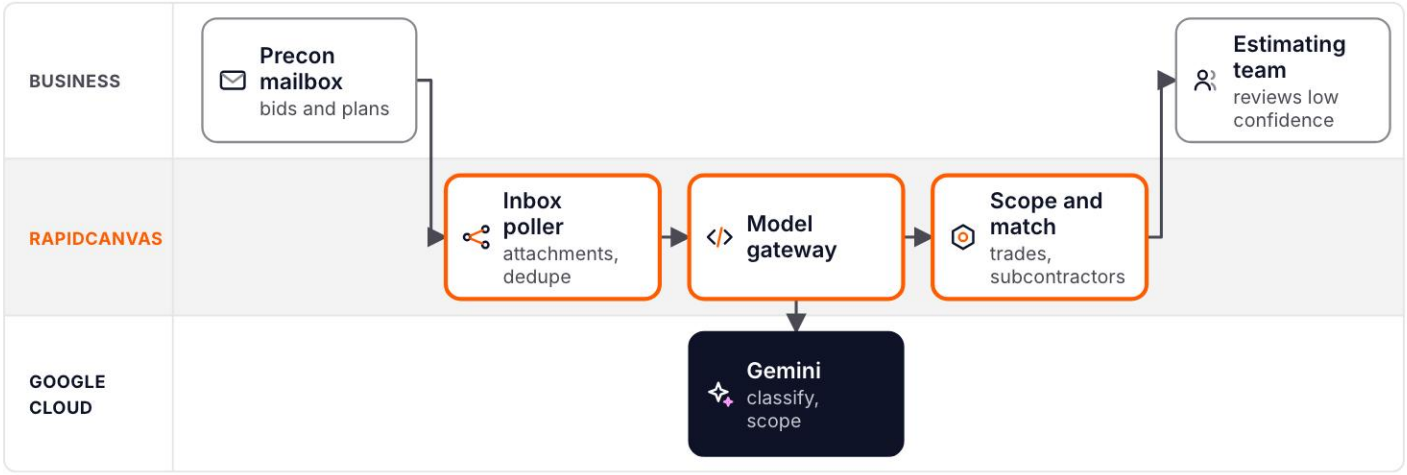
Gemini | AI Agents | Model gateway | Expert-in-the-Loop

The Challenge

A general contractor's estimating team triages a shared mailbox where subcontractor quotes, its own outgoing requests and marketing arrive together. Each real bid must be tied to the right project and subcontractor, every project scoped into divisions and trades from drawings, and bids evaluated line by line against that scope. Doing this by hand delays estimates and loses bids.

Solution Implemented

An inbox poller stores attachments, with rules resolving tagged replies first; everything else goes to Gemini via the gateway as a two-axis classification, applied only above 0.8 confidence. Gemini returns divisions filtered against a catalogue for uploaded drawings, forcing low-confidence picks to miscellaneous. Bid replies are evaluated by Gemini into line items, and the bidder is matched by email or a Gemini name match. A MySQL layer is wired for required TLS but stays inactive until credentials arrive.



How It Works

1. Poll the shared mailbox, store attachments in an artifact, and resolve tagged replies with rules.
2. Classify the rest with Gemini on two axes, applying results only above 0.8 confidence.
3. Scope each project into catalogue-constrained divisions and trades from uploaded drawings with Gemini.
4. Evaluate bid line items against the scope with Gemini, and match bidders by email or name.

Results and Benefits

- **A proof of concept since 2026, with the model gateway migration validated from inside the tenant.**
- **Low-confidence bid, project and subcontractor matches always route to a person, never guessed.**

HOW IT RUNS ON GOOGLE CLOUD

Gemini classifies, scopes and evaluates every bid on the same gateway that also reaches other models. The pattern is reusable regardless of the customer's own database.

USE CASE 12 | MULTIPLE INDUSTRIES, GLOBAL ● REFERENCE PATTERN

Structured Extraction and Managed Databases, Repeated

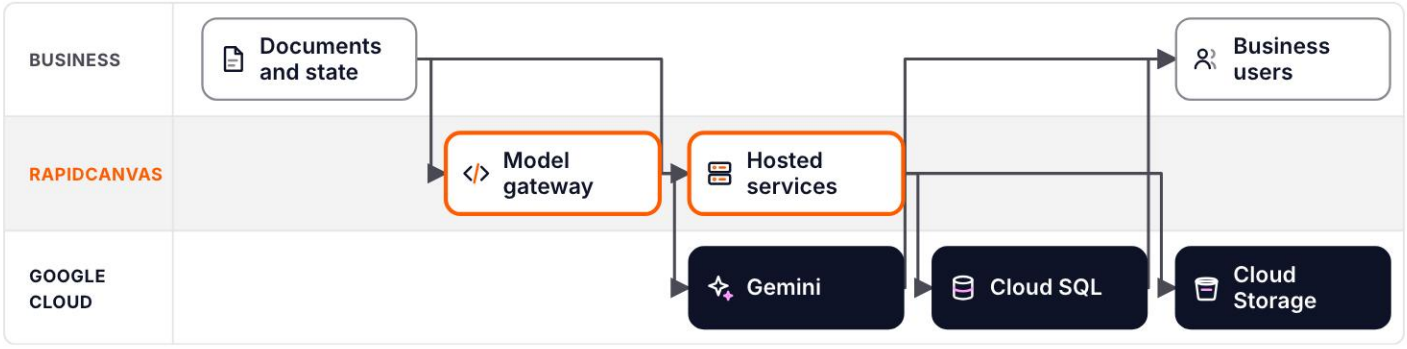
Gemini | Cloud SQL | Cloud Storage | Canvas pipelines | Model gateway | Connectors | Expert-in-the-Loop

The Challenge

Across many RapidCanvas deliveries on Google Cloud, two needs repeat outside any single flagship story. Teams must read a document or email and return a strictly shaped JSON object, validated and costed, with the right model picked per task. And hosted services need durable operational state in a managed database the customer already governs, without leaking credentials.

Solution Implemented

Seven further projects run Gemini behind the same validation layer: a fast model gates documents before a slower extractor reads them, structured e-mail parsing fills a strict schema, an agent answers questions with file-search tools, and a low-cost model narrates thousands of records after a cost comparison. Separately, at least eight deliveries keep operational tables in Cloud SQL for MySQL or PostgreSQL, reached over TLS or the Cloud SQL connector, with pool pre-ping and environment separation by database name.



How It Works

1. Gate documents with a fast Gemini model before an expensive extractor runs, where volume justifies it.
2. Constrain every Gemini call to a JSON schema, and record usage and cost for each one.
3. Host recipe and service state in Cloud SQL for MySQL or PostgreSQL, with pooled connections.
4. Separate environments by database name with fail-loud guards against production fallbacks.

Results and Benefits

- Three of seven Gemini-pattern projects are live, three in ramp or validation, one retired.
- Cloud SQL is confirmed behind at least eight RapidCanvas deliveries on Google Cloud.
- Batching one evaluation workload raised throughput by roughly two orders of magnitude.

HOW IT RUNS ON GOOGLE CLOUD

Every one of these smaller deployments still runs Gemini calls and Cloud SQL or Cloud Storage traffic. The breadth shows the pattern, not just flagship stories, drives consumption.

WORKING TOGETHER

RapidCanvas on Google Cloud

Every engagement runs on the Hybrid Approach™: the RapidCanvas agentic AI platform builds the pipelines, agents and Data Apps, and RapidCanvas experts work alongside the customer's team on data modeling, validation and the last mile into production.

RAPIDCANVAS BENCHMARKS ACROSS ENGAGEMENTS

4-8 weeks
to first production value

10x
faster time to value than a traditional custom build

Up to 80%
lower cost than a traditional custom build

Deploy where the data is

RapidCanvas runs as SaaS or inside the customer's own Google Cloud project, reading BigQuery, Cloud SQL, Cloud Storage, Drive and Sheets through IAM-scoped service accounts, with no data leaving the project by default.

Connect natively

Cloud Storage is a first-class RapidCanvas connector with browse, import and export. BigQuery, Cloud SQL, Drive and Sheets are reached through Google client libraries, with service-account keys held in RapidCanvas Secrets, never in code.

Route models through one gateway

Gemini Flash and Pro models run through the RapidCanvas model gateway, alongside other providers, with schema validation, rate limiting, retries and per-call token and cost accounting on every request.

Work with Google Cloud

RapidCanvas is a Google Cloud Partner listed in the Partner Directory, engaging customers directly today. A Google Cloud Marketplace listing is on the roadmap for teams that procure that way.

Customer names are withheld. Each story describes a real RapidCanvas deployment, verified against the delivered code. Delivery status is shown as of September 2026.

RapidCanvas

NEXT STEP

See How RapidCanvas Works on Google Cloud

See how RapidCanvas works on Google Cloud. Book a complimentary thirty minute discovery call.

Book a Discovery Call

rapidcanvas.ai



Real AI transformation, accelerated.