RapidCanvas

PARTNER CASE STUDY | RAPIDCANVAS + AWS

Twelve AWS Workloads, Built Inside Your Account

From contact-center analytics to fund-report extraction, RapidCanvas builds production AI on Amazon S3, Amazon Athena, Amazon Redshift, Amazon Connect and Amazon Bedrock, inside the customer's own AWS account.

12
use cases

7
in production

September 2026
edition



Real AI transformation, accelerated.

AT A GLANCE

RapidCanvas + AWS

RapidCanvas is a Managed AI Execution Company that ships production systems inside the infrastructure customers run. Across twelve AWS engagements, from contact-center quality analytics to private-markets fund-report extraction, RapidCanvas's Hybrid Approach™, an agentic platform paired with expert-in-the-loop review, attaches to Amazon S3, Amazon Athena, Amazon Redshift, Amazon Connect and Amazon Bedrock without moving data first.

Canvas pipelines, Data Apps, AI Agents and Ask AI run inside that account, and Amazon Bedrock stands alongside the direct Anthropic API as a configurable model provider. Seven of twelve use cases are already in production, from a distribution command center served from Amazon S3 to invoice reconciliation writing entries back into it. The rest move from pilot and proof of concept: RapidCanvas builds the AI and workflow layer, AWS keeps the data and compute.

12

AWS integration use cases documented

7

use cases already in production

11

distinct AWS services in active use

PARTNERSHIP

RapidCanvas is listed on AWS Marketplace and is an APN Partner, procured through private offers against committed AWS spend.

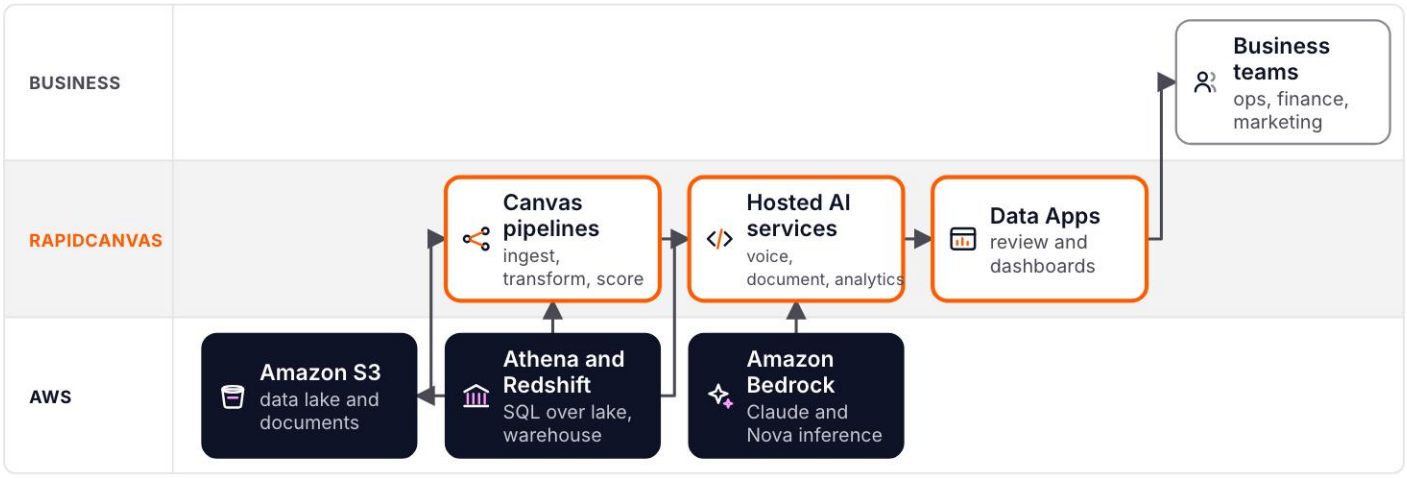
When RapidCanvas Is a Fit

1. The customer has data in Amazon S3, Athena or Redshift but no AI team.
2. A contact-center or document workload needs Bedrock or Transcribe inference with a human reviewer.
3. Compliance rules mean the workload must run entirely inside the customer's own AWS account.
4. The customer wants Bedrock evaluated against a direct model API before choosing one provider.

REFERENCE ARCHITECTURE

How RapidCanvas Builds on AWS

The same pattern recurs across all twelve engagements: credentials from IAM, files through Amazon S3, analytics through Athena and Redshift, inference through Bedrock, RapidCanvas doing the AI and workflow layer in between.



Read data in place

RapidCanvas connects to Amazon S3, Amazon Athena, Amazon Redshift, Amazon Connect and Amazon Bedrock with scoped IAM credentials, and never asks a customer to copy data into a separate platform-owned store first.

Deploy inside the account

RapidCanvas runs as SaaS or as a private deployment on Amazon EKS provisioned by Terraform inside the customer's own AWS account, inheriting that account's IAM, KMS and network posture.

Keep the model configurable

Amazon Bedrock stands alongside the direct Anthropic API as a model backend selected by configuration, so inference can move inside the AWS boundary without a rebuild of the application.

AWS Services in These Stories

Amazon S3
Landing zone and durable store for documents and data lakes

Amazon Bedrock
Configurable Claude and Nova inference alongside the direct Anthropic API

Amazon Athena
SQL queries over S3 data lakes and Parquet extracts

Amazon Redshift
Warehouse source of truth for marketing and finance marts

Amazon Connect
Contact-center trace records and call recordings for nightly analytics

Amazon Transcribe
Real-time streaming speech to text for a voice agent

Amazon RDS
Relational store for claims, features and test results

AWS KMS
Customer-managed encryption keys for data, database and secrets

AWS Secrets Manager
Runtime secrets and credentials for containerized services

Amazon ECS on AWS Fargate
Containerized compute for a platform deployed inside one AWS account

Amazon DynamoDB
Low-latency configuration and conversation state for a voice agent

IN THIS DOCUMENT

12 Use Cases, 11 AWS Services

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

NO.	USE CASE	INDUSTRY	STATUS
01	From Thousands of Recordings to Nightly Scorecards	Contact center operations, Latin America	**In production**
02	Turning a Claims Backlog Into a Review Queue	B2B marketing technology, global	**In production**
03	Seven Dashboards Consolidated Into One Live View	Energy distribution, Americas	**In production**
04	Media Kits and Clicks Become One Recommendation Engine	Affiliate marketing, global	**In production**
05	Quarterly Fund Reports Turned Into Structured Data	Financial services, global	**In production**
06	A Governed Lakehouse Replaces a Sixty-Table Legacy Schema	Real estate, Latin America	**In production**
07	Commission Invoices Reconciled Straight From S3 Datamarts	Affiliate marketing, Europe	**In production**
08	One Reconciled Dashboard Replaces a Dozen Spreadsheets	Consumer skincare, North America	**In delivery**
09	Supplier Spreadsheets Map Themselves to the Schema	Retail, Europe	**In delivery**
10	A Voice Agent Routes Sales Calls in Real Time	Telecommunications, Latin America	**Reference pattern**
11	A Test-Data Platform Deployed Inside One AWS Account	Advanced manufacturing, North America	**Proof of concept**
12	Interview Transcripts Become a Searchable, Grounded Chat	Management consulting, global	**Proof of concept**

USE CASE 01 | CONTACT CENTER OPERATIONS, LATIN AMERICA ● IN PRODUCTION

From Thousands of Recordings to Nightly Scorecards

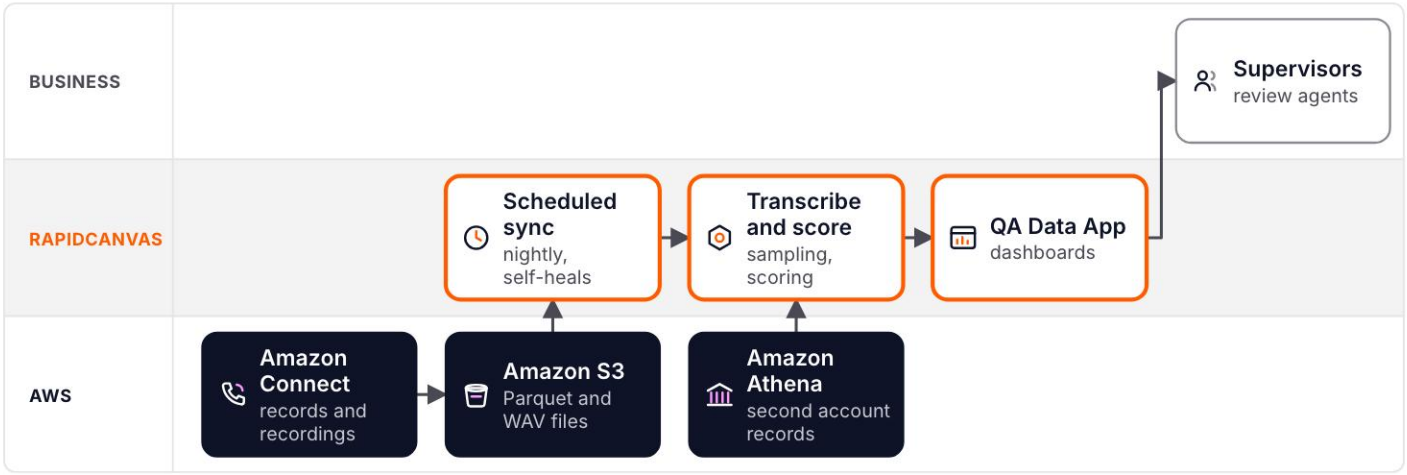
Amazon Connect | Amazon S3 | Amazon Athena | Scheduled jobs | Data Apps | Connectors

The Challenge

A contact center for an energy distributor in Latin America needed continuous quality review of agents on clarity, courtesy, objectivity, knowledge and resolution, plus sentiment and call-type mix, without supervisors listening to thousands of recordings. Records ran to millions of rows, and a second channel's calls sat behind Amazon Athena in a separate account.

Solution Implemented

Amazon Connect exports contact trace records and recordings as Parquet and WAV to Amazon S3 daily. A scheduler wakes before midnight, pulls voicebot records from the second account through Amazon Athena, and syncs Amazon S3 recordings into an analytics store. Calls over fifty seconds are sampled at ten percent a day, water-filled per agent, transcribed, and scored against five criteria and six call types. Results feed a Data App with quality, agent and sentiment dashboards.



How It Works

1. Amazon Connect exports contact trace records and call recordings to Amazon S3 every day.
2. A scheduler wakes before midnight and self-heals with a catch-up run if a night is missed.
3. Athena supplies the second account's voicebot records; S3 supplies human-agent recordings and metadata.
4. Ten percent of eligible calls are transcribed and scored nightly across five agent criteria.

Results and Benefits

- In production, with millions of contact records and hundreds of thousands of calls analysed to date.
- Ten percent of eligible calls are scored every night across a bank of transcription workers.
- Supervisors review agents from dashboards instead of listening to recordings.

HOW IT RUNS ON AWS

Amazon Connect and Amazon S3 hold every contact record the nightly analysis reads, and Amazon Athena serves the second account's data in place, so consumption grows with call volume.

USE CASE 02 | B2B MARKETING TECHNOLOGY, GLOBAL ● IN PRODUCTION

Turning a Claims Backlog Into a Review Queue

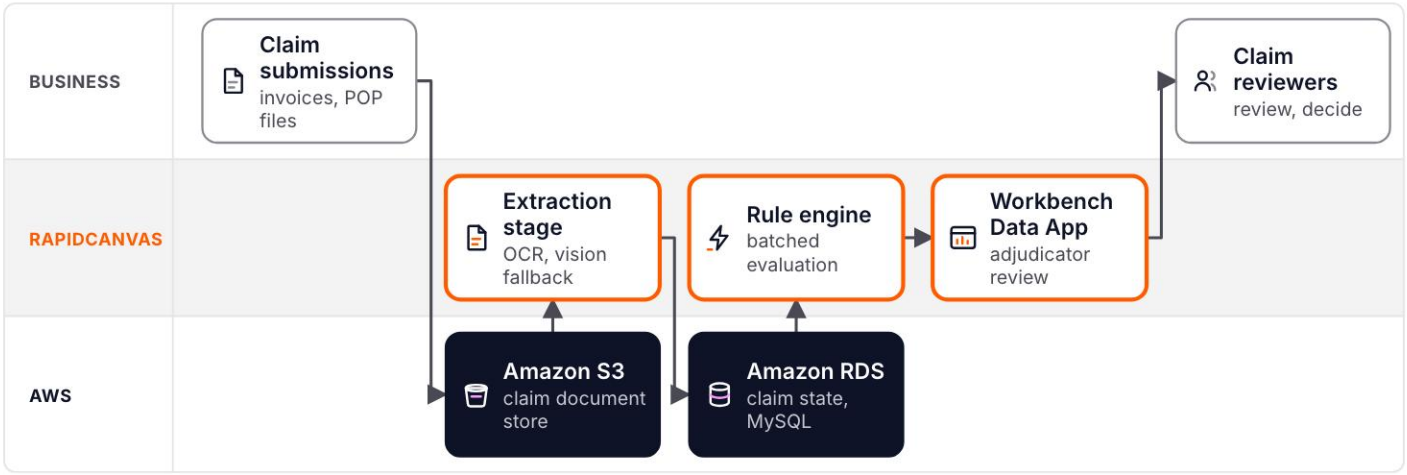
Amazon S3 | Amazon RDS | Document extraction | Expert-in-the-Loop | Data Apps | Connectors

The Challenge

A B2B marketing technology provider validates dealer reimbursement claims against roughly five hundred policy rules covering document presence, totals, rates, caps and creative compliance. Most invoices and proof-of-performance files arrive as an unknown type, many scanned, at a backlog of tens of thousands a day, before any rule could run.

Solution Implemented

Claim documents are uploaded to Amazon S3, and claim metadata syncs into a MySQL database on Amazon RDS. An extraction stage pulls each file from Amazon S3, routing small images and PDFs to vision OCR and larger PDFs to text extraction with a vision fallback. One prompt classifies each document and returns a class-specific schema. A rule engine evaluates the five hundred rules, batching model checks per document, and adjudicators finish review in a workbench Data App.



How It Works

1. Claim documents are uploaded to Amazon S3, the only durable store on the serving platform.
2. An extraction stage pulls files from Amazon S3 and routes them by type and size.
3. One prompt classifies each document and returns a class-specific structured schema.
4. A rule engine batches checks per document, cutting model calls by about ninety percent.

Results and Benefits

- **One programme is already live; a second is in validation.**
- **Nearly all validation-set documents processed successfully with model analysis.**
- **The rule base of about five hundred rules is validated against thousands of claims.**

HOW IT RUNS ON AWS

Amazon S3 is the only durable store for every claim document the models and rule engine touch, and Amazon RDS holds claim state, so consumption scales directly with claim and document volume.

USE CASE 03 | ENERGY DISTRIBUTION, AMERICAS ● IN PRODUCTION

Seven Dashboards Consolidated Into One Live View

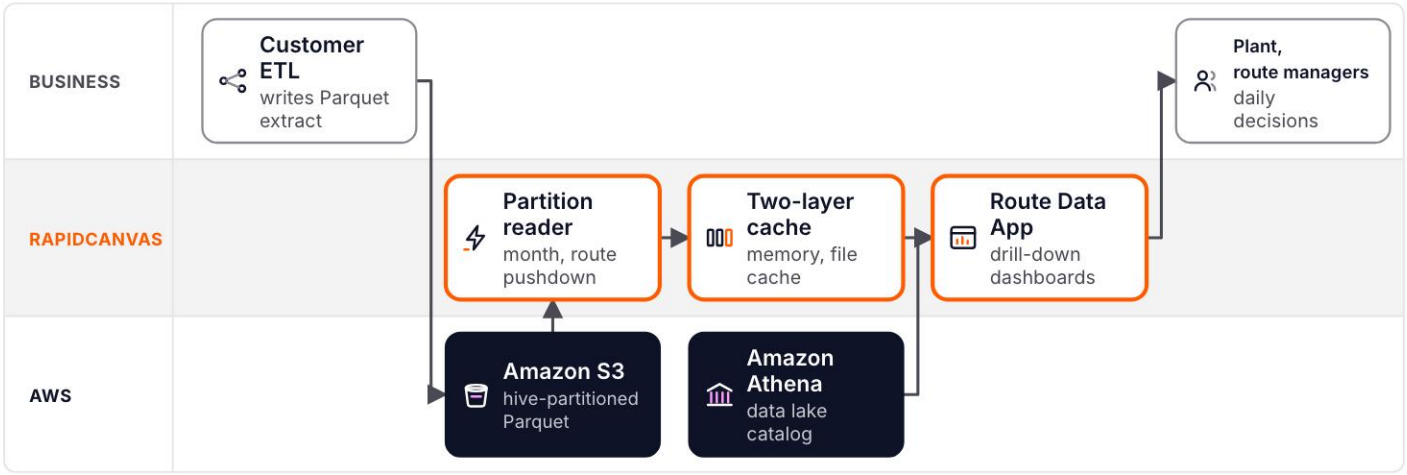
Amazon S3 | Amazon Athena | Data Apps | Connectors | Canvas pipelines

The Challenge

Plant and route managers at an energy distributor needed one view of daily service orders across several legal entities, plus street sales. The data existed as Parquet in the company's AWS lake, but the view was spread across seven microservices, with a thirty-second budget per request.

Solution Implemented

The customer's ETL writes a services extract to Amazon S3 as Parquet, catalogued for Amazon Athena. A hosted service consolidates seven former APIs, reading only the month partitions a request needs from Amazon S3, then filtering to exact dates in memory. Street sales are served from the same data platform. A two-layer cache keeps repeat queries inside the time budget, and a Data App drills from date to plants, routes and customers with cascading filters.



How It Works

1. The customer's ETL writes a Parquet services extract to Amazon S3, catalogued for Amazon Athena.
2. One hosted service replaces seven microservices, reading only the month partitions a request needs.
3. Predicate pushdown on month and route keeps single-customer drill-downs from scanning a full month.
4. A two-layer cache keeps repeat queries, including street sales, inside the time budget.

Results and Benefits

- Already live, serving month-to-date queries over hundreds of thousands of rows.
- Every request completes within the thirty-second worker budget from partition-pruned reads.
- Managers work from one drill-down Data App instead of seven separate services.

HOW IT RUNS ON AWS

Amazon S3 stays the customer's data lake and Amazon Athena the query engine over it; RapidCanvas only prunes partitions and caches results, growing AWS usage with the business.

USE CASE 04 | AFFILIATE MARKETING, GLOBAL ● IN PRODUCTION

Media Kits and Clicks Become One Recommendation Engine

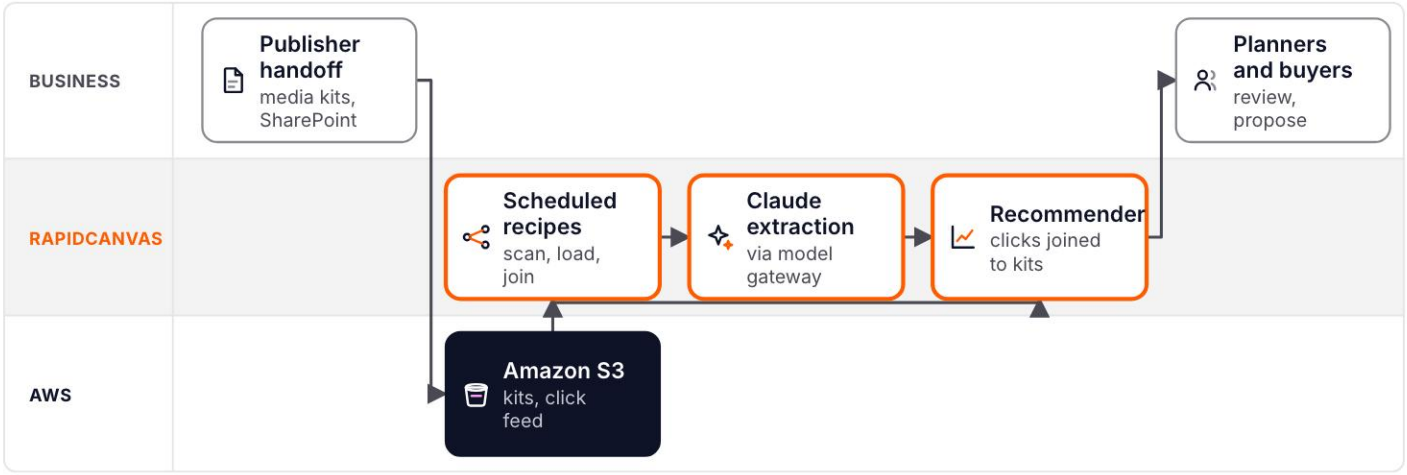
Amazon S3 | Document extraction | Model gateway | Expert-in-the-Loop | Data Apps

The Challenge

An affiliate-marketing network's planners buy publisher placements from about two hundred media kits a quarter, arriving as PDFs in varying formats. Evaluating a placement relied on memory, each proposal took thirty to forty-five minutes, and the network's click-performance history, hundreds of millions of rows, already lived as Delta and Parquet in Amazon S3.

Solution Implemented

The customer's IT team copies publisher PDFs into a dated Amazon S3 prefix daily. A scheduled recipe resolves the publisher and quarantines unresolved kits. An extraction library reads each media-kit family with Claude through the model gateway, emitting structured data with per-field provenance. A second recipe loads the performance feed's Delta and Parquet files from Amazon S3 into a clicks table. Derived tables join clicks to placements, and a recommendation engine produces proposals buyers finalise from a review queue.



How It Works

1. The customer copies publisher PDFs into a dated Amazon S3 prefix every day.
2. A scheduled recipe resolves the publisher, de-duplicates kits and quarantines anything unresolved.
3. Claude, called through the model gateway, extracts structured fields with per-field provenance.
4. A second recipe loads the Amazon S3 click-performance feed into a clicks fact table.

Results and Benefits

- Already live, with a further milestone in ramp.
- The clicks table holds hundreds of millions of rows in each environment.
- A dry run processed hundreds of media-kit PDFs from the Amazon S3 test prefix.

HOW IT RUNS ON AWS

Amazon S3 is both the media-kit handoff and the hundreds-of-millions-of-rows click-performance store the recommendation engine reads every day, so the AWS footprint grows with placement volume.

USE CASE 05 | FINANCIAL SERVICES, GLOBAL ● IN PRODUCTION

Quarterly Fund Reports Turned Into Structured Data

Amazon Bedrock | Data Apps | Document extraction | Model gateway | Expert-in-the-Loop

The Challenge

Analysts at an outsourced financial-services provider turn quarterly PDF reports from private-credit fund managers into a per-instrument dataset with dozens of fields. Reports differ in layout and vocabulary, feed investor reporting, so extraction had to be accurate and cheap at scale, with some clients needing inference inside their AWS boundary.

Solution Implemented

An analyst uploads a PDF to a Data App, which picks the best input format per page. An orchestrator splits the document into page ranges and enforces extract-only rules. The model provider is chosen by configuration: a Bedrock prefix routes to Amazon Bedrock through the Converse API for Claude, while an Anthropic prefix calls the API directly. Derived metrics are computed in Python, never the model, and analysts review, override and export results to Excel.



How It Works

1. An analyst uploads a fund PDF; the app picks the best input format per page.
2. An orchestrator splits the document and enforces extract-only rules on every field.
3. Configuration routes the same request to Amazon Bedrock or the direct Anthropic API.
4. Python derives metrics like loan-to-value and TVPI, never the model, then analysts review.

Results and Benefits

- **Already live alongside a second use case for the same customer.**
- **Bedrock support is production code, selected per deployment by configuration.**
- **Analysts review, override and export every extraction to Excel or CSV.**

HOW IT RUNS ON AWS

Amazon Bedrock gives the customer a configuration-only path to keep Claude inference inside its own AWS account, so the same extraction application can move from direct-API to AWS-resident inference without a rebuild.

USE CASE 06 | REAL ESTATE, LATIN AMERICA ● IN PRODUCTION

A Governed Lakehouse Replaces a Sixty-Table Legacy Schema

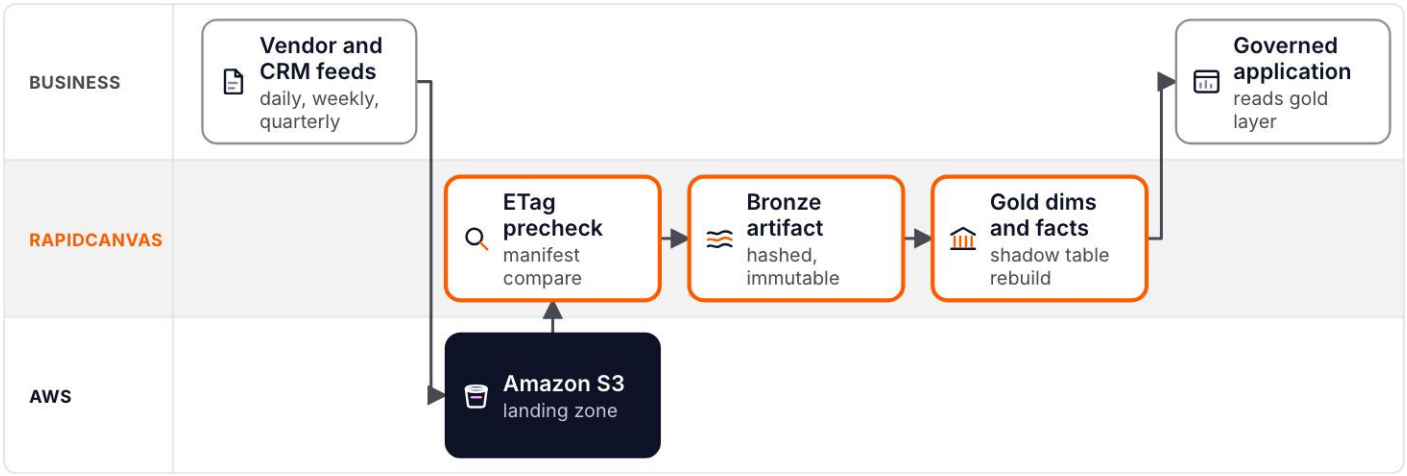
Amazon S3 | Canvas pipelines | Scheduled jobs | Connectors

The Challenge

A residential developer's market-intelligence system consolidates its inventory and sales with CRM, market data and public statistics. Vendor files land in Amazon S3 on varying cadences, often unchanged, and silent schema drift kept breaking loads. A legacy schema of sixty tables mixed state with pipeline output, with no lineage.

Solution Implemented

Vendor CSVs and spreadsheets arrive under a root prefix in the customer's Amazon S3 bucket. A pre-check recipe compares each object's Amazon S3 ETag against a manifest, so unchanged files are skipped and contract failures block only that partition. Approved files are hashed and persisted to an immutable bronze artifact, then rebuilt into gold dimensions and facts through atomic table swaps, so the application never reads a half-built table. Data-quality rules run after every build.



How It Works

1. Vendor files land in Amazon S3 under a dated prefix on their own cadence.
2. Amazon S3 ETags are compared to a manifest so unchanged files are skipped.
3. Approved files are hashed and persisted to an immutable bronze artifact.
4. Gold dimensions and facts rebuild through shadow tables and atomic renames.

Results and Benefits

- **Already live, with more than a dozen ingest recipes running daily.**
- **Dozens of pre-checks and gold builders run without manual intervention.**
- **The application now reads one documented gold layer instead of an ad hoc schema.**

HOW IT RUNS ON AWS

Amazon S3 is both the vendor landing zone and the content-addressed bronze store behind every recipe, so as vendor feeds and file volume grow, so does the customer's own S3 consumption.

USE CASE 07 | AFFILIATE MARKETING, EUROPE ● IN PRODUCTION

Commission Invoices Reconciled Straight From S3 Datamarts

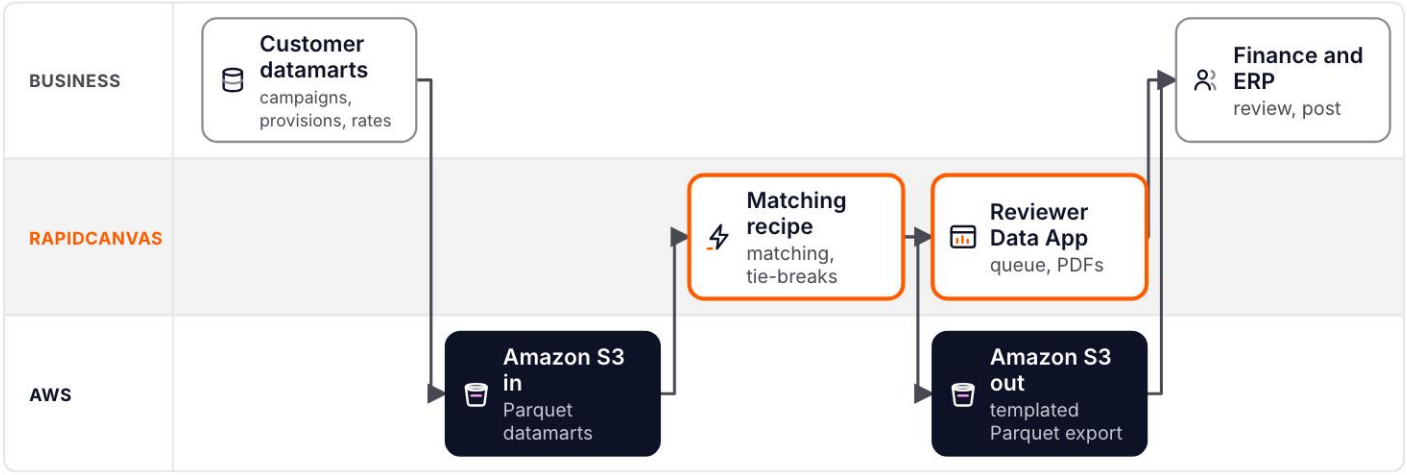
Amazon S3 | Canvas pipelines | Document extraction | Data Apps | Connectors

The Challenge

Finance at a European affiliate-marketing group receives commission invoices from dozens of networks, in many currencies, mostly PDFs. Each invoice must match campaigns already sold, convert at the right rate, check against provisions and post to the ERP. Datamarts publish this data as dated Parquet in Amazon S3.

Solution Implemented

The customer's datamarts publish dated Parquet folders for campaigns and rates to an Amazon S3 input bucket. A daily recipe lists each prefix, picks the newest folder and reads the Parquet in chunks. API-connected invoices are downloaded and their descriptions extracted through the model gateway. A reconciliation recipe matches invoice lines to campaigns with documented tie-breaks and converts amounts per month. The final export writes back to Amazon S3 as Parquet for the ERP, reviewed in a Data App.



How It Works

1. Datamarts publish dated Parquet folders for campaigns and rates into Amazon S3.
2. A daily recipe reads the newest Amazon S3 folder in chunks.
3. A reconciliation recipe matches invoice lines to campaigns with documented tie-breaks.
4. The templated export writes back to Amazon S3 as Parquet for the ERP.

Results and Benefits

- Already live for API-connected invoices; PDF ingestion is in validation.
- Most payments matched the ERP's postings in a recent check.
- Invoices verified duplicate-free across two years of history.

HOW IT RUNS ON AWS

Amazon S3 carries both the datamarts RapidCanvas reads and the accounting entries it writes back, so every invoice adds to the customer's own S3 usage.

USE CASE 08 | CONSUMER SKINCARE, NORTH AMERICA ● IN DELIVERY

One Reconciled Dashboard Replaces a Dozen Spreadsheets

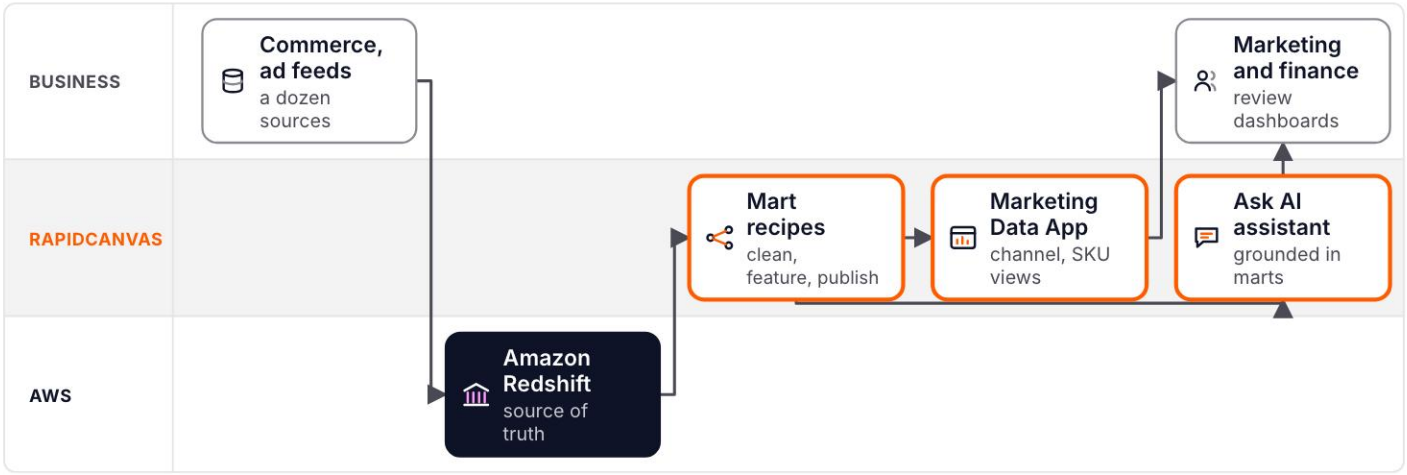
Amazon Redshift | Canvas pipelines | Data Apps | Ask AI | Connectors

The Challenge

A direct-to-consumer skincare brand already landed revenue and spend from a dozen sources in Amazon Redshift, but marketing and finance read them through dashboard sheets whose definitions differed by tab. They wanted one reconciled view of channel performance, acquisition cost and return on ad spend, tied out against incumbent reports.

Solution Implemented

Amazon Redshift stays the source of truth. Ingest recipes query its modelling schemas read-only on a rolling 180-day window, and feature recipes build channel, campaign, retention and SKU marts. An outputs recipe publishes each mart with an atomic rename, avoiding partial refreshes. A hosted Data App serves channel, campaign and SKU views, and Ask AI answers from the same marts. Every mart was reconciled sheet by sheet against the incumbent dashboard.



How It Works

1. Four recipes read Amazon Redshift's modelling schemas over a read-only connection daily.
2. Clean and feature recipes merge feeds into channel, campaign, retention and SKU marts.
3. An outputs recipe publishes each mart with an atomic rename, never a partial refresh.
4. Every mart is reconciled sheet by sheet against the incumbent dashboard reports.

Results and Benefits

- **Channel revenue closely ties out to the incumbent reports on a settled window.**
- **Dozens of typed marts replaced flat CSV outputs, with matching indexes.**
- **Several metric-dictionary fixes were found and corrected during reconciliation.**

HOW IT RUNS ON AWS

Amazon Redshift remains the customer's warehouse and source of truth; RapidCanvas only reads it on a rolling window, so warehouse usage keeps growing with the business.

USE CASE 09 | RETAIL, EUROPE ● IN DELIVERY

Supplier Spreadsheets Map Themselves to the Schema

Amazon Bedrock | Data Apps | Model gateway | Connectors | Expert-in-the-Loop

The Challenge

A European retail group onboards supplier spreadsheets in German and English into a product-information system. Column names vary by supplier and language, with unit columns and barcode synonyms, and an engineer hand-mapped every column each shipment, with confirmed mappings never reused. The customer wanted ranked, explained candidates with model-provider choice, including AWS-resident inference.

Solution Implemented

A user drops a supplier workbook into a Data App, which loads the target schema and history from an artifact store. A similarity pre-filter shrinks each batch of columns to a short list before any model call. A dispatcher routes by prefix: Bedrock-prefixed requests use Amazon Bedrock with IAM keys or an Identity Center token; other prefixes call Anthropic, OpenAI or Gemini directly. The model returns three ranked, explained candidates per column, saved for reuse.



How It Works

1. A supplier workbook is parsed and matched against the target schema and history.
2. A similarity pre-filter narrows each batch of columns before any model call runs.
3. Bedrock-prefixed requests use Amazon Bedrock; other prefixes call APIs directly.
4. The model returns three ranked, explained candidates per column for the user to confirm.

Results and Benefits

- **Runs on a dedicated tenant for the customer and is in validation.**
- **Bedrock is a first-class provider selected by configuration, alongside the direct API.**
- **Confirmed mappings and prompt edits are reused automatically on the next shipment.**

HOW IT RUNS ON AWS

Amazon Bedrock is selected purely by configuration, so the customer can move Claude inference into its own AWS account for this workload without changing the application.

USE CASE 10 | TELECOMMUNICATIONS, LATIN AMERICA ● REFERENCE PATTERN

A Voice Agent Routes Sales Calls in Real Time

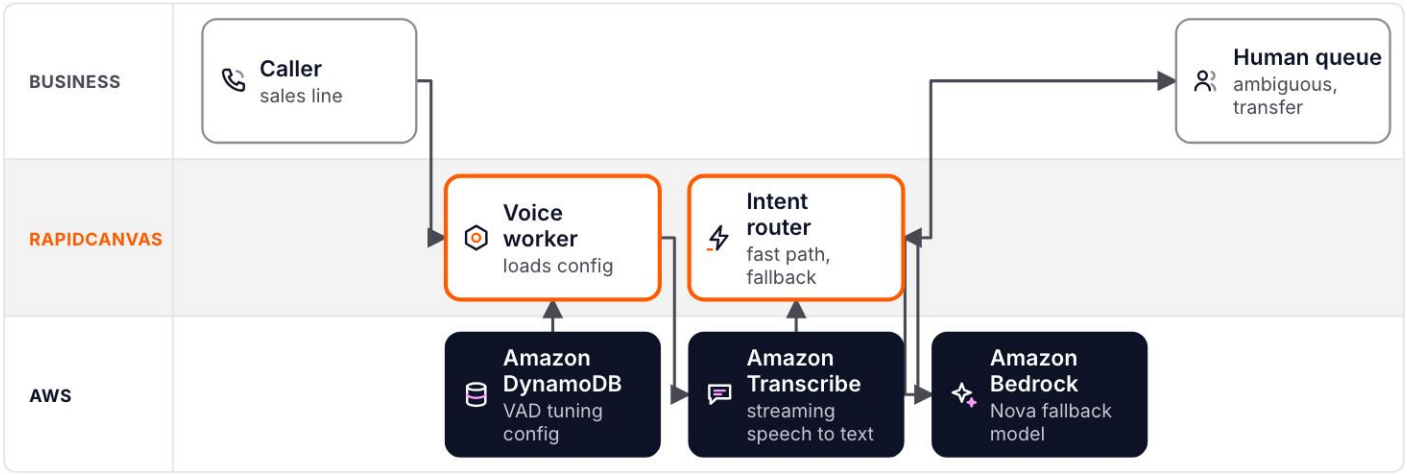
Amazon Transcribe | Amazon Bedrock | Amazon DynamoDB | AI Agents | Model gateway | Connectors

The Challenge

A telecom operator's sales line receives calls with mixed intents: faults, billing, cancellations and mobile-plan questions, alongside genuine purchase interest, and agents lose time on callers who will never buy. The first deliverable was narrow: greet the caller, understand intent within a turn or two, and route to the right human queue quickly.

Solution Implemented

A RapidCanvas voice worker loads tuning settings from Amazon DynamoDB, so tuning never needs a redeploy. Every turn streams to Amazon Transcribe as audio; only final transcripts pass on, and filler noise is dropped. The router resolves intent with a lexical fast path first and falls back to a Bedrock Nova model through the Amazon Bedrock Converse API, with two ambiguous turns routed to a human advisor. Cached, pre-approved responses return audio in milliseconds.



How It Works

1. A voice worker loads tuning settings from Amazon DynamoDB for every call.
2. Each turn streams to Amazon Transcribe; only final, filtered transcripts move on.
3. A fast path resolves intent first, falling back to a Bedrock Nova model.
4. Two ambiguous turns in a row route the caller straight to a human advisor.

Results and Benefits

- **The classifier resolves tested intents at or near 100% on the measured corpus.**
- **Bedrock classification takes well under a second per fallback turn.**
- **A reference build: the production phone trunk and live transfer are not yet connected.**

HOW IT RUNS ON AWS

Amazon Transcribe and Amazon Bedrock carry every call turn the agent handles, so contact-center volume converts directly into transcription and inference consumption if the pattern moves toward production.

USE CASE 11 | ADVANCED MANUFACTURING, NORTH AMERICA ● PROOF OF CONCEPT

A Test-Data Platform Deployed Inside One AWS Account

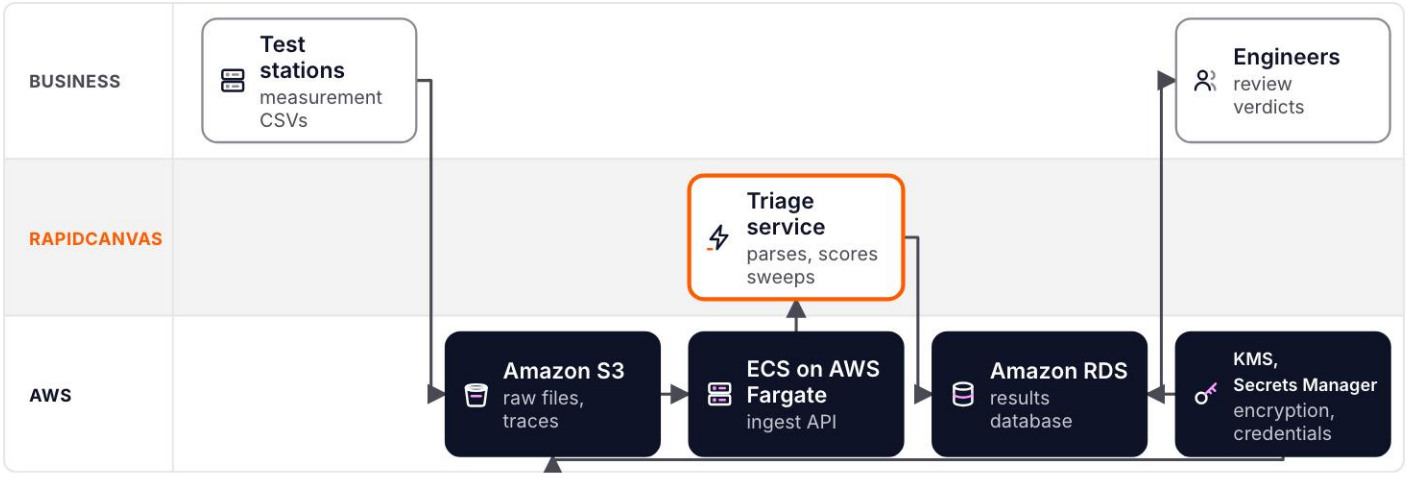
Amazon ECS on AWS Fargate | Amazon S3 | Amazon RDS | AWS KMS | AWS Secrets Manager | Data Apps | Connectors | Expert-in-the-Loop

The Challenge

A regulated hardware manufacturer runs stations that export per-unit measurement sweeps as thousands of CSV files per unit, but the machine verdict is a blunt pass or fail, so engineers spend time separating genuine failures from benign notches. Because the data is controlled it cannot leave the account, ruling out a shared SaaS tenant.

Solution Implemented

The whole platform is defined in Terraform for the customer's AWS account: Amazon ECS on AWS Fargate for compute, Amazon RDS for results, Amazon S3 for raw files and traces, AWS KMS for encryption and AWS Secrets Manager for credentials. The same stack deploys to a commercial or government-cloud region. Engineers upload files or point the service at ones already in Amazon S3, encrypted under one key. A background job triages each file for engineers to confirm.



How It Works

1. Terraform deploys one container service, a database and storage into the customer's account.
2. Engineers bulk-upload files, or point the service at files already in Amazon S3.
3. One customer-managed AWS KMS key encrypts storage, database, secrets and logs alike.
4. A background job triages sweeps; engineers confirm the verdict on a review screen.

Results and Benefits

- **A proof of concept already running inside the customer's own AWS account, encrypted under one key.**
- **Engineers triage thousands of sweep files per unit from one review screen.**

HOW IT RUNS ON AWS

Amazon ECS on AWS Fargate, Amazon S3, Amazon RDS, AWS KMS and AWS Secrets Manager all run inside the customer's AWS account, targeting commercial or government-cloud regions.

USE CASE 12 | MANAGEMENT CONSULTING, GLOBAL ● PROOF OF CONCEPT

Interview Transcripts Become a Searchable, Grounded Chat

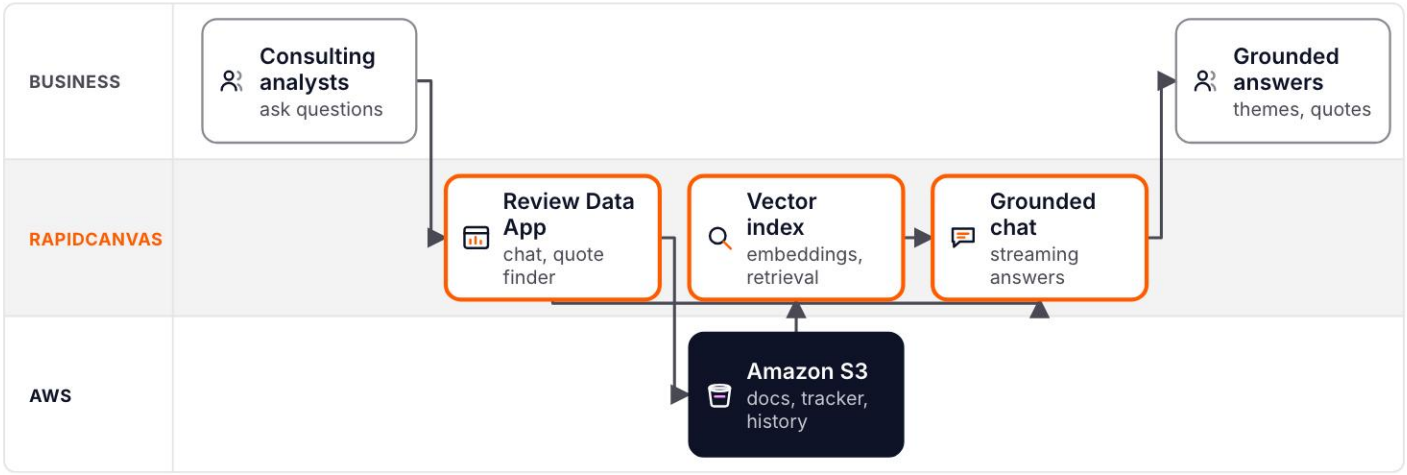
Amazon S3 | Ask AI | Data Apps | Connectors

The Challenge

Analysts at a management-consulting firm read dozens of interview transcripts per engagement to find themes and quotes. The firm wanted to ask questions across every project document with grounded answers, filter by client segment, locate exact quotes and pull in reference material, without standing up a new document system.

Solution Implemented

Each project is a prefix in the firm's own Amazon S3 bucket; a tracker file, also in Amazon S3, records each project. Transcripts and references upload to their prefix, and the service streams them in for parsing. Documents are chunked and embedded for a quote finder, and a tag selector infers client segments from the question. Chat history and feedback round-trip a workbook, also kept in Amazon S3, so deleting a project removes everything at once.



How It Works

1. Every project is one prefix in the firm's own Amazon S3 bucket.
2. Transcripts and references upload to their prefix and stream in for parsing.
3. Documents are chunked and embedded so a quote finder can ground every answer.
4. Chat history and feedback round-trip a workbook that also lives in Amazon S3.

Results and Benefits

- Already in use with the firm as a proof of concept.
- Amazon S3 provides the entire storage layer, so transcripts never leave the firm's account.
- Analysts ask questions across every transcript instead of reading them one by one.

HOW IT RUNS ON AWS

Amazon S3 is the only storage layer this engine touches, for documents, the vector tracker and chat history, so every new project adds directly to the firm's S3 usage.

WORKING TOGETHER

RapidCanvas on AWS

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 as a private deployment on Amazon EKS, provisioned with Terraform: VPC, managed node groups, Amazon RDS databases and load balancers stood up inside the customer's own AWS account.

Connect natively

Amazon S3 and Amazon Redshift are first-class RapidCanvas connectors. Amazon Athena, Amazon DynamoDB, Amazon Transcribe and Amazon Connect data are reached through the AWS SDK, with credentials held in RapidCanvas Secrets.

Choose the model backend

Amazon Bedrock is a supported model backend, Claude, Amazon Nova, Amazon Titan, Mistral and Llama models, alongside the direct Anthropic API. Model choice is configuration, and every call carries usage accounting and lineage.

Buy through AWS Marketplace

RapidCanvas is listed on AWS Marketplace and is an APN Partner, procured through private offers, with EDP eligibility subject to the customer's own AWS contract.

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 AWS

See how RapidCanvas runs inside an AWS account, then book a discovery call to scope your first use case.

Book a Discovery Call

rapidcanvas.ai



Real AI transformation, accelerated.