AdRevora AI — Comprehensive Technical Overview

Current release and source publication — 14 October 2026

This current section supersedes earlier current-state descriptions below. The previous overview remains intact as a dated record. AdRevora AI and its AdRevora OS workspace are deployed with the frontend, authenticated revenue backend and private CPU inference on the existing server. Release 20261014-platform-07 is verified; the inference release remains unchanged. The full ten-phase program continues through company activation, pilot observation, expansion and distributed acceptance.

Delivered reporting and financial controls

The latest release adds deterministic CEO, marketing and operations metrics, inspectable immutable creative revisions with exact approval history, and retained daily UTC aggregate reports. Completed-day captures replay the original snapshot; the scheduled command runs after a successful encrypted backup. New summaries are protected by the next backup. CSV includes campaign, unknown-attribution and separately marked total rows with signed numeric money and reporting metadata. Campaign plus unknown money reconciles to the total; contact counts can overlap campaigns and must not be summed as money.

Exact supplied-record provenance prevents one fresh cost delta from refreshing unrelated CRM or finance evidence. Legacy provenance remains unknown, replays do not refresh evidence, and conservative freshness covers retained source records. Stable bounded history pages reuse calculations by query, while configured creation/journal ceilings prevent unbounded work. The reference ledger remains $9,500 net cash with zero reconciliation difference; these maintained reference records describe financial behavior, not company outcomes.

Verification, deployment and recovery

Fresh publication gates passed 120 application tests across 14 files, 56 deployment tests, types, lint, formatting and build; the dependency advisory query reported zero vulnerabilities. Five security gates passed, including public/history secret scans, Bandit, actual Semgrep coverage and the installed edit scanner. Independent review approved the source and public guides. The existing approximately 516 KB frontend chunk advisory remains documented.

Real HTTPS checks passed 17 workflow checks and 23 scoped accessibility scans with zero reported application errors or violations, plus 21 static and three API checks, certificate validation and HTTP-to-HTTPS redirect. Flows included password/TOTP, actual private CPU inference at zero vendor API cost, reconciled CSV, persona metrics, stored daily reports, exact revision decisions and stale-approval rejection. Exact viewport widths passed at 320, 390, 768 and 1440 pixels. Full manual accessibility acceptance remains open.

Native inspection matched 68 release artifacts, four installed units, proxy, private environment and runtime; 12 other site configurations were unchanged. All 20 table hashes were preserved across the final upgrade; scheduled daily generation replayed without duplicates. Measured core interruption was 5.288 seconds. A 295,011-byte encrypted backup restored in isolation with eight key table hashes matching, zero restored sessions, plaintext removed and an off-server ciphertext hash match. These observations do not establish a monthly availability guarantee or disaster-recovery objective.

GitHub publication and complete public guides

Only after a normal source push was this existing native overview inspected. GitHub main commit 995b2c1deedaa3c86d57e2b783ed207d2a40ba1a was verified exactly; all 149 source file identities and all five README contents matched. GitHub associates author and committer with Zahidulislam2222. Repository visibility remains private and Actions remains disabled. About, homepage and project topics were updated and read back. Source delivery and deployed artifact parity were checked separately; this documentation revision required no new application promotion.

The repository guides cover product and architecture, frontend/backend setup, API and data model, Meta/GoHighLevel/Stripe activation, Attribution and CAPI, AI authority, legal/privacy, threat controls, accessibility, operations, backup/restore and rollback, ownership, costs and phased delivery. The new release decision register assigns accountable functions and evidence to each future gate. Documentation navigation checks passed 49 guides and 274 local links. Source: https://github.com/Zahidulislam2222/AdRevora-AI . Running workspace: https://adrevora.zahidul-islam.com/engine/ .

Legal, security and future scale acceptance

The 14 October primary-source review covers GDPR and EU tracking/transfer duties, California privacy and automated-decision requirements, advertising claims, communications, payment-data boundaries and sector applicability. Future health-related processing requires separate HIPAA/FTC rule review. Qualified legal/privacy owners approve applicability, consent, rights, retention, transfers and incident decisions. Broader security acceptance maps versioned ASVS 5.0.0 controls and requires independent assessment; future PostgreSQL roles must account for owner, superuser and BYPASSRLS exceptions. Software tests alone do not provide those approvals.

The future workload is 1.1M simultaneous active dashboard sessions plus a separate 1.1M tracked visitors. At the stated activity assumptions this means 110,000 API requests/second and about 36,667 events/second. Continuous tracking would yield 3.168 billion events/day, about 3.168 decimal TB/day at 1 KB/event before replicas and overhead. The roadmap migrates reconciled transactions to tenant-constrained PostgreSQL, durable queues/outbox, separate ingestion/reporting/provider/AI workers, materialized analytics, scoped caches and multi-AZ regional tenant cells. Acceptance requires representative tenant skew, failure headroom, cold/warm caches, stepped load, a two-hour peak, 24-hour soak, recovery results and measured unit costs.

The planning availability floor is 99%; the future monthly objective is 99.9% for agreed console/API journeys and separately for durable event acceptance. Thirty-day time budgets are 432 and 43.2 minutes; request-based error budgets use their own denominator. Define eligible journeys, dependency failures and maintenance before an SLA. External probes, burn alerts, on-call ownership, canaries and tested rollback support the objective. Proposed core RPO of five minutes and RTO of sixty minutes require measured drills. Phase 10 follows accepted company integration, legal/security review, the 30-day pilot, business expansion and capacity/availability evidence. No new paid service, provider write or remote CI was activated by this publication.

Previously retained overview — earlier dated scope

Reviewed 14 October 2026 | Latest verified product milestone: 13 October 2026

Latest review — release evidence and reading map

Reviewed 14 October 2026. This expanded overview adds implementation and handover detail to the existing concise overview, which follows in full. The latest verified product milestone remains the 13 October release. The document review did not rerun production tests or establish a newer application release. Earlier recorded results below retain their original scope and date.

The delivered product combines a public multi-page website with an authenticated revenue workspace and optional private CPU inference on the owner's existing server. The accepted single-server release includes persistent finance and attribution, mandatory MFA, versioned creative approvals, team membership, reports and exports, consent-aware conversion records and encrypted recovery. Company account activation remains a separate next step.

The release record reports 107 application tests and 54 deployment tests, with types, lint, formatting, build and security checks passing. Sixteen HTTPS browser checks covered authentication, financial evidence, approval behavior, exports, conversions, connector readiness and actual private inference. These results support the scoped implementation; they do not establish real customer adoption, current company-source completeness or sustained service availability.

Upgrade evidence records 63 matching release files plus installed service and runtime configuration, exact value preservation in all 20 production database tables, and twelve other hosted sites unchanged. The measured core interruption was 3.077 seconds. Actual browser inference completed in 34.868 seconds with zero vendor API cost for that run. Server hosting and future operations still require their own budget and ownership.

Reading map: the next six sections explain everyday workflows, implemented architecture, financial semantics, bounded AI, recovery and launch acceptance. The retained overview then provides the complete earlier concise narrative, ten-phase summary, capacity arithmetic and historical milestones. Native headings make these sections navigable through the document outline.

Evidence hierarchy: the private dossier supplies verified narrative; the maintained source guides define detailed contracts; dated verification records describe what actually passed. Product proposals remain labeled as proposals. The Bennett brief supplies requirements rather than proof of a client engagement, and the maintained revenue scenario is an arithmetic reference rather than a business outcome.

The practical next milestone is an authorized, controlled company integration and pilot with agreed finance definitions and accountable owners. It requires account permissions, source mappings, completeness and reconciliation proof, privacy/security approval and support arrangements. Publishing this document does not activate providers, enable campaign writes or authorize spending.

User workflows — evidence, decisions and collaboration

Access begins with an operator-provisioned identity and authenticator enrollment through an approved private channel. The hosted workspace requires a password and a current authenticator code; password-only access is rejected. An authenticated session exposes only workspaces covered by current membership. Users should sign out on shared devices and handle downloaded reports under the organization's retention rules.

Before reading a campaign result, choose the reporting period, attribution model and campaign. Check the source mode, currency, freshness, completeness and unknown attribution. Reference workspaces use maintained records. An imported workspace reflects only the source snapshots accepted by its ledger. The interface cannot turn an incomplete source period into complete business history.

Administrators can create recommendations and record approve or reject decisions. Analysts can inspect authorized reports, journeys, current drafts and decision history. CEO, marketing and operations report perspectives change guidance labels; they do not create extra authentication roles or distinct permission-scoped datasets. Administrator and analyst membership remains the actual authorization boundary.

A recommendation is reviewed against its exact report query, evidence and policy version. A changed evidence snapshot or stale recommendation requires another review; stale mutations can return HTTP 409. Approval records a human decision in the workspace. It does not publish an advertisement, change an account budget or execute a provider action in the current release.

Creative work is internal and revisioned. An administrator saves a meaningful brief, inspects the resulting revision and approves or rejects that revision. An edit creates a new revision and invalidates the previous approval. The current listing shows the latest draft and its decision state; retained storage does not imply an unrestricted revision-history API.

Administrators manage memberships for existing users. Revocation invalidates affected sessions, and the last administrator is protected against accidental removal. Public registration and enterprise identity provisioning remain separate roadmap items. Exports enforce current workspace access and report policy for each request, with CSV formula prefixes escaped to reduce spreadsheet execution hazards.

Connector readiness is inspected before synchronization. Company Meta and HighLevel access, approved mappings and real bounded validation are prerequisites. First-party conversions require stable occurrence identity and valid permission. Privacy suppression cancels pending marketing delivery while required financial evidence is retained. External conversion delivery remains disabled pending activation and acceptance.

Implemented architecture — ingestion to trusted reports

The frontend uses React, TypeScript and Vite, with sixteen route entries covering the public product experience and authenticated engine. Per-route HTML, self-hosted assets and reduced-motion behavior support delivery. A same-origin Node HTTP service owns authenticated operations and SQLite persistence. This is the implemented single-server baseline; the proposed distributed platform has a separate migration and acceptance contract.

The request path is browser over HTTPS, session verification, current workspace membership, validated operation, ledger or workflow service, and an authorized response. Signed provider webhooks have a separate ingress boundary based on raw-body signature, account, event mode and version checks. Browser session rules cannot substitute for provider verification, and a caller-supplied workspace identifier cannot grant membership.

Storage separates identities and access, canonical records and imports, connector ownership and runs, recommendations and decisions, creative revisions, conversion occurrences, delivery state, suppression and audit. Foreign keys and strict tables support relationships. Nested domain payloads are validated by application code; SQL typing alone cannot validate every monetary field or authorize a tenant operation.

Source ingestion validates identity, monetary precision, timestamps, relationships and control totals before transactional acceptance. A failed bounded synchronization preserves the last successful source snapshot. Meta reads account and campaign information plus daily spend; HighLevel supplies approved consent, touch, stage and enrollment evidence. Account bindings prevent silently replacing one company's populated workspace with another account.

Journey attribution connects approved touches and enrollment evidence to authoritative money without rewriting cash totals. Reports show known and unknown attribution explicitly. Source freshness and completeness accompany the result. Source identifiers can still identify people through linkage even when names, email and phone are discarded; they remain protected data rather than automatically anonymous records.

One typed settings boundary owns environment and deployment values. Maintained mappings, prompts and policy belong in validated data files; actual credentials remain outside public source and documents. Provider origins, API versions, model identity, limits and timeouts are configuration. Stable protocol enums remain distinct from changeable business rules. A configuration change still needs the applicable validation and review.

Distributed migration would add reviewed PostgreSQL tenant constraints and row-level security, durable jobs and an outbox, stateless identity, analytical read models and admission control. It needs schema migration, source and finance parity, bounded backfill and rollback evidence. The current reviewed upgrader rejects schema changes; the future architecture must be implemented and measured before capacity claims are made.

Financial meaning — cash, attribution and reconciliation

Successful collections and successful refunds are recorded separately using validated integer minor units and an allowed currency. Net collected cash is collections minus refunds. An enrollment stage, opportunity value or advertising platform conversion claim is evidence about the journey; it does not establish that cash was received. Missing acquisition or fulfillment costs remain explicit rather than being assumed zero.

The maintained reference scenario records $10,000 collected and $500 refunded, producing $9,500 net. Attribution assigns $7,000 to known evidence and leaves $2,500 unknown. Known plus unknown equals net cash, so the reconciliation difference is zero. This is supplied-record arithmetic. It does not prove actual revenue, live provider completeness or full accounting profit.

The proposed education journey starts with a Meta ad and webinar registration, continues through a HighLevel contact, Platinum appointment, completed consultation and enrollment, then connects verified payment to a recommendation. The supplied offer price is $5,000. Discounts, installments and refunds can make actual cash differ from that nominal price, so reports require agreed source definitions and enrollment identity.

Delivery identity and financial-object identity are separate. Two authentic webhook events may describe one unchanged payment; accepting both deliveries must not duplicate money. A changed amount, currency or enrollment relationship must be rejected rather than treated as another harmless replay. Signed ingestion covers supported successful payment and refund events; a real company receiver and replay exercise remain prerequisites for activation.

Attribution requires approved touch evidence, reporting boundaries and valid consent. Unknown or missing relationships stay visible. A report end date is a filter and must not silently age earlier in-window leads out of the result. Human decisions bind the selected reporting period and evidence, preventing an approval for one campaign period from becoming blanket approval for a different period.

Before finance use, compare a bounded source period with company provider exports and payment records. Check counts, pagination, time zones, revised daily spend, duplicated or reordered events, refunds and missing relationships. Repeating a synchronization should retain stable totals while accepted source revisions remain journaled. Establish a historical completeness strategy and an owner for unresolved exceptions.

Fees, chargebacks, taxes, overhead and multiple currencies require deliberate accounting extensions and acceptance. Recorded contribution depends on known costs and is distinct from full business profit. Financial retention and legal holds also differ from marketing-data deletion. Restoration must preserve money and reapply current access and privacy decisions rather than blindly reviving an older operating state.

Private AI — bounded authority and observed behavior

The optional private model supports investigation of campaign evidence. The server sends numeric aggregates and policy labels rather than raw contact or campaign identifiers. The model selects allowlisted investigation codes; deterministic code calculates financial facts and renders their explanation. Extra prose, unknown codes and repeated codes are rejected. The model cannot rewrite source records, money, permissions or approvals.

The accepted runtime uses pinned Ollama 0.40.3 and Qwen2.5 1.5B identities with verified model content. The recorded manifest and five configuration or layer blobs total 986,061,405 bytes. A pinned Nginx 1.30.4 gateway exposes only required methods and routes. These are dated deployed-version facts, not recommendations to adopt those versions without current maintenance review.

The model runs under a dedicated nonroot identity with two CPU cores, 3 GiB RAM, 128 processes, a read-only root, bounded temporary storage and logs, one inference worker, a two-request model queue and a 2,048-token context. Its owned internal network has no published port or default route. Cloud features were disabled and CPU-only warmup was observed during acceptance.

The separate gateway has one CPU core, 128 MiB RAM and 32-process bounds, with nonroot and read-only controls. It provides the sole loopback publication and has a fixed upstream destination. Its edge network has a different egress boundary from the model's internal network. It receives no application credentials or database, and administrative or model-download routes are denied.

Native checks covered ten route, method and encoded-path cases; attempts to redirect the upstream through Host or query; refusal through the server's nonloopback address; and actual private-address change followed by DNS recovery. Model outage and restart were exercised, with core readiness maintained. Recovery verified ownership and restored the exact runtime rather than relaxing isolation guards.

Actual application HTTP checks used an isolated in-memory reference workspace. Controlled timeout cancellation occurred after 1,008 milliseconds with peer socket closure observed. Concurrent analysis returned HTTP 429 while financial reads took 4–6 milliseconds and password login 899–1,926 milliseconds during pending analysis. These are recorded bounded observations, not service-level latency commitments.

The isolated inference completed in 39.808 seconds; a later real HTTPS browser inference completed in 34.868 seconds. Both recorded zero vendor API cost, and the supplied finance remained unchanged. AI unavailability is explicit, leaving policy and evidence available for human review. There is no paid cloud fallback in this runtime; future model changes require evaluation, licensing, operational review and applicable authorization.

Security and recovery — preserving current authority

Implemented access controls include password hashing, bounded login attempts, encrypted authenticator secrets, replay counters, expiring hashed sessions and Secure, HttpOnly, Strict cookies. Each authorized operation checks current workspace membership. Exact host and origin checks constrain browser mutations. Provider webhook ingress instead checks its signed raw payload and bound account, mode and version.

Negative verification covers unauthorized roles and tenants, stale approvals, duplicate or ambiguous finance, hostile origins, unsafe model output and formula-bearing CSV values. Audit records retain workspace, actor, action and time. These controls and scanner results support the current bounded release; independent penetration testing, threat acceptance and broader identity lifecycle review remain launch gates.

Recovery uses an encrypted consistent SQLite snapshot, an isolated restore and an encrypted off-server copy. Recorded checks covered database integrity, foreign keys and current finance, access, MFA and suppression values; restored sessions were cleared. A successful isolated drill proves that bounded exercise. Daily backups alone do not establish the future five-minute recovery point objective.

Access and consent cannot be treated as ordinary historical business data. Restoring an older database must not revive revoked membership, accept an already-used authenticator code or reinstate withdrawn marketing permission. Current authority and consent histories are reconciled, suppression tombstones reapplied and affected queued delivery canceled. Financial reconciliation remains an explicit required step alongside restore integrity.

The accepted upgrade took an encrypted safety backup before promotion and preserved values in all 20 tables. Local-to-live checks verified release and configuration parity. Earlier contained deployment failures led to real nonroot filesystem validation, private validator storage and socket-bind checks before mutation. Inference corrections added native restricted gateway validation and ownership-bound recovery.

Operations need named primary and backup maintainers, incident command, escalation, monitoring and reviewed rollback. Monitor errors, source freshness, unknown cash, queue state, database pressure and model availability. A provider outage should retain the last trustworthy evidence with its freshness status rather than fabricating current data. Record restored authority and source reconciliation during every material recovery exercise.

Before real company onboarding, accountable owners must decide lawful purposes, roles, retention, withdrawal, rights requests, international transfers and vendor terms. Raw card data stays outside this analytics service. Sector expansion and AI/media use need their own review. The public overview excludes credentials and private operational detail; controlled recovery records remain separate from shareable documentation.

Next acceptance — company pilot and scalable release

The immediate activation gate is company ownership and permission. Confirm the intended Meta account, HighLevel location, payment account and workspace binding through approved access. Verify currently supported API versions and minimum scopes when activation occurs. Supply approved CRM mappings, source timestamps, consent history and enrollment identifiers; missing business evidence cannot be invented by the connector.

Run a bounded company synchronization and compare source counts and daily money with authoritative exports. Verify pagination, account currency and time zone, historical coverage, revisions, replay behavior and failure handling. Register the dedicated payment receiver with its own signing secret and test signed payment and refund cases. Tracking requires asset audit, approved event mapping, consent, deduplication and company Test Events acceptance.

A controlled pilot needs product acceptance, finance reconciliation, privacy/legal sign-off, independent security review and named support owners. Define the user cohort, reporting periods, success measures, freshness expectations, incident route and rollback before onboarding. Evaluate recommendations against known evidence and mature outcomes; local reference arithmetic and one successful inference are not pilot business results.

The ten-phase plan covers interface, secure foundation, provider connections, tracking, journey and revenue, AI and approvals, initial pilot, wider integrations, scale and reliability, and final handover. Implemented pieces span several phases, while each phase retains explicit exits. Google and wider business journeys require separate permissions and definitions. Proposed provider writes need human authority, budget safeguards and an explicit activation decision.

The future capacity model separates 1.1 million dashboard sessions at one request per ten seconds from 1.1 million tracked visitors at one event per thirty seconds: 110,000 API requests and about 36,667 events per second before caching or batching. The current server has not proven that throughput. Regional cells, queues, analytical reads and tenant admission need measured load, failure, backlog and cost evidence.

Future objectives are a 99.9% monthly availability target above a 99% floor, RPO no greater than five minutes and RTO no greater than sixty minutes for agreed data classes. A 30-day 99.9% time-based allowance is 43.2 minutes. Define eligible journeys and dependencies before commitments; sustained external measurement and actual recovery exercises are required to accept those targets.

Handover should identify the accepted source and build, configuration version, rollback artifact, operating owners, source coverage and unresolved findings. Account and infrastructure costs require a verified budget and separate authorization. The following retained overview preserves all earlier current-document wording, including its history summary and source guide map, so this expansion supplements rather than replaces the prior content.

Current position

AdRevora AI is an advertising-to-revenue intelligence product; AdRevora OS is the platform name. The public website and authenticated revenue backend run over HTTPS on the owner's existing server. The current release includes private, bounded CPU inference. The Bennett brief is a requirements reference, not evidence of a client engagement. The supplied financial example is a maintained reference scenario, not customer revenue.

The 13 October release met its scoped existing-server delivery and AI-upgrade acceptance contracts. It does not complete the proposed ten-phase program. Company-owned Meta, HighLevel and Google accounts, a dedicated Stripe receiver, provider completeness checks, independent legal and security approval, a controlled pilot, distributed capacity and measured monthly availability remain open gates.

Purpose and first business journey

The product links advertising spend and creative evidence to CRM stages, appointments, enrollment, collected cash and explainable recommendations. The first planned connected journey is Meta ad → webinar registration → lead → HighLevel contact → Platinum appointment → completed consultation → enrollment → verified revenue → recommendation. The supplied offer price is $5,000, but actual cash depends on payments, discounts and refunds. The main proposed KPI is cost per Platinum enrollment; revenue per lead is secondary.

The data model keeps unknown attribution explicit. A source record is not silently treated as a sale, and collected cash is not automatically called profit. Later proposed journeys include real estate transactions, property acquisition and insurance, each requiring its own definitions, permissions and compliance review.

Delivered product and user workflows

The React, TypeScript and Vite website has sixteen route entries, including product, solutions, integrations, security, resources, legal pages and the authenticated engine. Per-route HTML, self-hosted visual assets and reduced-motion behavior support the public experience. The backend is a Node HTTP service with SQLite persistence on the current single server.

An authorized user can sign in with a password and mandatory replay-resistant MFA; open a tenant workspace; inspect campaign and journey evidence; review reconciled finance and attribution; request a bounded recommendation; create and approve a versioned creative draft; manage scoped team membership; export role-limited CSV or JSON; record consent-aware conversions; and sign out. Editing an approved draft invalidates that approval. Changes to membership revoke affected sessions and protect the last administrator.

Provider delivery remains disabled pending company activation. Meta and HighLevel adapters are read-only and were exercised against controlled fixtures and HTTP contracts. Stripe ingestion verifies signed, account-, mode- and version-bound events, but real company receiver registration and reconciliation remain pending. The product does not autonomously publish campaigns or spend advertising funds.

Architecture and trusted data

The existing release uses a same-origin frontend and API. A typed settings boundary owns deployment values; maintained mappings, prompts and policy live in data files. The backend persists

tenant access, source snapshots, journey evidence, money postings, attribution, decisions, drafts, consent and audit records. Financial writes have a single controlled path with replay and duplicate protections. Exports enforce tenant role and purpose boundaries.

Source snapshots retain provenance and freshness. The ledger records collection and refunds separately; attribution connects money to evidence without changing the authoritative cash total. The reference scenario has $10,000 collected and $500 refunded, leaving $9,500 net. Of that, $7,000 is attributed and $2,500 remains unknown; the reconciliation difference is zero. These figures verify arithmetic on supplied records. They do not establish live provider completeness or a customer result.

Consent and suppression are checked before future marketing delivery. Withdrawal cancels pending delivery while retaining required financial evidence. The local service does not claim network-wide exactly-once delivery. Historical imports, out-of-order events, source gaps, chargebacks, fees, taxes and multiple currencies need explicit accounting rules and acceptance before finance use expands.

AI and human authority

The recommendation path uses numeric aggregates and policy labels, not raw contact or campaign identifiers, as model input. The model selects from allowlisted investigations; server code calculates financial facts and applies policy. AI cannot alter cash, approval authority, source state or provider settings. A human remains responsible for decisions and any later activation of external writes.

The current optional model runs on the existing server under a dedicated nonroot identity with bounded CPU, memory, process count, queue and context. The model network has no published port or default route. A fixed-route gateway is the only loopback publication and holds no application credentials or data store. Native checks covered model-byte integrity, isolation, DNS recovery, outage and restart, timeout cancellation and overload rejection. In an actual HTTPS browser flow, private inference completed in 34.868 seconds with zero vendor API cost. This is an observed run, not a latency guarantee.

Security, privacy and operating boundaries

Current controls include password hashing; expiring Secure, HttpOnly, Strict sessions; encrypted MFA secrets; explicit tenant roles; input and operation bounds; exact host and origin checks; signed webhook validation; audit records; and protected exports. The release also exercised encrypted backups, an isolated restore and an off-server encrypted copy. The restored database passed integrity and foreign-key checks, with finance, access, MFA and suppression values preserved and restored sessions cleared.

Before real company data is onboarded, accountable owners must settle controller and processor roles, lawful purposes, consent and withdrawal, retention and deletion, cross-border transfers, vendor terms and sector-specific advertising rules. Independent threat review, penetration testing, rights-request exercises and launch-blocking finding closure remain required. Raw card data is outside this analytics service. Future AI and media use needs classification, disclosure and license review. These are acceptance gates, not certifications already earned.

Production operations need named primary and backup maintainers, on-call coverage, incident command, monitoring for errors and source freshness, reviewed rollback and measured restore exercises. A successful backup and one isolated restore do not prove the future recovery objectives.

Verification of the current release

The scoped 13 October acceptance recorded 107 application tests and 54 deployment tests passing, plus type, lint, formatting, build and security gates. Semgrep examined 108 actual files with zero findings or errors. Sixteen live browser checks at mobile through desktop widths reported zero application errors; the flow included MFA, finance, approvals, role exports, conversion records, pending connectors, logout and actual private inference. HTTP and TLS checks also passed. Independent implementation and deployment reviews accepted the scoped release.

The local-to-live release comparison covered 63 files and installed runtime configuration. All 20 production database tables retained identical values during the AI upgrade, and twelve other hosted sites were unchanged. Measured core interruption was 3.077 seconds. The source and final documentation revisions were pushed normally with verified owner attribution and exact remote commit equality; the repository remains private and remote Actions disabled. These observations establish this release's bounded state, not future uptime, provider coverage or million-user capacity.

Ten-phase plan and remaining exits

1. Frontend design and product interface — delivered and published; further product acceptance remains per phase.

2. Secure foundation — authenticated persistent single-server base delivered; enterprise identity and final distributed architecture review remain.

3. Meta and HighLevel connections — read-only adapters and fixtures delivered; company account authorization, mappings, bounded sync and completeness remain.

4. Tracking — first-party conversion records, consent and suppression delivered; asset audit, Pixel/CAPI mapping and company Test Events remain.

5. Journey and revenue — local ledger, evidence, attribution and export flows delivered; real finance reconciliation and expanded accounting remain.

6. AI and approvals — bounded private inference, human decisions and creative draft approvals delivered; product evaluation and approved provider actions remain.

7. Initial connected pilot — requires company activation, privacy/security sign-off, controlled use, support and measured outcomes.

8. Wider integrations — Google/GA4 and other businesses require separate permissions, data definitions and acceptance.

9. Scale and reliability — distributed implementation, funded load and failure tests, recovery proof and measured objectives remain.

10. Final release and handover — requires completion and acceptance of every applicable product, legal, operational and contractual gate.

The first proposed connected release spans phases 1–7; shipping parts of those phases does not close the full phase exits. Cloud and SaaS choices remain proposals until reviewed and approved.

Capacity, availability and cost boundaries

The future workload contract models 1.1 million dashboard sessions at one request per ten seconds, plus 1.1 million tracked visitors at one event per thirty seconds. That is 110,000 API requests per second and about 36,667 events per second before caching or batching. It is planning arithmetic, not proven throughput. The proposed architecture separates static delivery, authorized APIs, event

collection, durable queues, transactional money, analytical reads and bounded AI jobs; tenant admission limits and regional failure headroom need actual tests.

The objective is 99.9% monthly availability for defined critical console/API and event-acceptance journeys, above a 99% floor. In a 30-day month, 99.9% allows 43.2 minutes of time-based unavailability. The proposed recovery goals are RPO no more than five minutes and RTO no more than sixty minutes. None of these future targets has been accepted through sustained measurement. Capacity purchases, paid APIs and cloud deployments require a separately approved budget.

Earlier milestones retained in summary

The initial design milestone produced the public multi-page website and responsive route coverage. Subsequent local core work added authenticated finance, attribution, connector contracts, deterministic policy and persisted human decisions. Local acceptance reached 14/14 criteria with 89 application tests, four deployment tests and thirteen browser checks before the hosted backend milestone.

The first hosted release added persistent MFA, team roles, versioned drafts, scoped reports and exports, consent-aware conversions, encrypted recovery and local/live parity. A later private CPU upgrade introduced isolated inference while preserving financial and access data. Earlier documentation editions recorded source-only delivery and a static frontend; those are historical states superseded by the verified 13 October release. Their full wording remains available in this Google Doc's version history and the owner's retained pre-condensation PDF.

The implementation log also records contained release failures and corrections: shared-parent permissions blocked nonroot startup; a validator needed private temporary storage; an occupied port required an explicit socket-bind gate; and the private inference network and gateway needed native runtime checks rather than configuration assumptions. Each correction was made locally, reviewed and verified before the accepted release.

Source guides and handover

The owner repository contains the build plan, technical overview, API reference, data model, connector setup, user guide, AI governance, security architecture, compliance, operations, scalability, deployment, verification and final delivery records. Those detailed guides own executable contracts and evidence; this overview is the short current-state map. Company accounts and licenses remain owner-controlled. A later public case study should use only verified outcomes and omit private operational details.