# AIM Previz

# Information Security Policy Pack & Evidence

Version 1.0 · July 2026

Nine core policies · endpoint security · architecture & data flow · six evidence documents

PROVIDED FOR ENTERPRISE SECURITY REVIEW

# AIM Previz — Information Security Policy Pack

**Version 1.0 · Adopted 2026-07-06 · Review cadence: quarterly, and on material change Owner: Security Owner (interim: Founder/CEO) · Applies to: AIM Previz platform, all personnel and contractors with production or customer-data access**

This pack is the policy set behind the public statement at aimpreviz.com/security. It maps 1:1 to the nine policy areas enterprise security reviews request when no third-party attestation exists yet, and to the SOC 2 readiness program in `docs/compliance/soc2-readiness.md` (these policies close that program's "Policies" gap list).

| # | Policy | File | Primary review questions served |
|---|---|---|---|
| 1 | Access Control | `access-control-policy.md` | IAM, least privilege, MFA, password policy |
| 2 | Data Protection, Retention & Deletion | `data-protection-policy.md` | Encryption, classification, deletion SLAs |
| 3 | Patch & Vulnerability Management | `patch-vulnerability-management-policy.md` | Scanning, severity SLAs, pen-test cadence |
| 4 | Secure Development (SDLC) | `secure-development-policy.md` | Change management, code review, release flow |
| 5 | Incident Response Plan | `incident-response-plan.md` | Detection, severity, customer notification |
| 6 | Security Awareness Training | `security-awareness-training-policy.md` | Personnel training cadence |
| 7 | Vendor Management & Subprocessor Register | `vendor-management-policy.md` | Upstream AI vendors, no-training terms |
| 8 | Business Continuity & Disaster Recovery | `business-continuity-dr.md` | Backups, RTO/RPO, restore testing |
| 9 | Logging & Monitoring | `logging-monitoring-policy.md` | Event capture, retention, alerting |
| — | Endpoint & Device Security | `endpoint-security-policy.md` | Managed-device standards, BYOD |
| — | Architecture & Data Flow | `architecture-data-flow.md` | System boundaries, GenAI entry/exit points |

**Honesty rule (inherited from the platform's public claims):** every statement in these policies describes either (a) a control that exists in the architecture today, or (b) a commitment with a named owner and cadence. Nothing is claimed as certified or audited until a CPA-issued report exists.

Distribution: shareable under NDA as part of enterprise security review.

# Access Control Policy

**v1.0 · Adopted 2026-07-06 · Owner: Security Owner (interim: Founder/CEO)**

## 1. Principles

- **Least privilege.** Access to systems that store or process customer data is granted only on documented business need, scoped to the minimum required.
- **Individual accountability.** Every user has a unique, individual account. Shared or generic accounts are prohibited for interactive login.
- **Database-enforced authorization.** Tenant isolation is not an application convention: every table carries row-level security (RLS) policies and every storage bucket uses tenant-/project-scoped paths. A request without the right identity cannot read another tenant's rows, regardless of application bugs.

## 2. Production access

- Production infrastructure (database, storage, edge functions, secrets) is accessible only to the Security Owner and explicitly authorized engineers.
- Service-role credentials exist only in server-side edge-function environments; they are never present in client code, repositories, or developer laptops.
- Model-vendor API keys are server-side only. Client applications have no direct path to any AI vendor.

## 3. Authentication requirements

- **MFA is required** on every system handling customer data or production configuration: Supabase, GitHub, Google Workspace, Stripe, and the domain registrar. Phishing-resistant factors (passkeys/FIDO2, or app-based TOTP where passkeys are unsupported) are the standard; SMS is not an acceptable factor.
- **Product accounts:** passwords are stored only as one-way cryptographic hashes (bcrypt via the auth layer); plaintext passwords are never stored or logged. The auth layer rate-limits failed sign-ins. Sessions use short-lived JWTs with automatic expiry and refresh; sessions are revocable server-side.
- **Enterprise engagements:** operator accounts serving an enterprise engagement enforce minimum 12-character passwords (14 for administrative roles) and MFA. SAML 2.0 / OIDC single sign-on is available on the dedicated-tenant deployment tier.

## 4. Access lifecycle

- **Grant:** requested with business justification, approved by the Security Owner, recorded (date, scope, approver).
- **Review:** all production and vendor-console access is reviewed **quarterly**; the review is recorded with reviewer and outcome.
- **Revoke:** on role change or departure, access is removed the **same day**; offboarding uses a written checklist covering all systems in §3.

## 5. Customer-triggered revocation

An enterprise customer may trigger full access revocation for its data at any time via its named contact channel. On trigger: engagement accounts are suspended, share links revoked, and signed URLs expire — target completion within **1 hour** of the request, 24/7.

# Data Protection, Retention & Deletion Policy

**v1.0 · Adopted 2026-07-06 · Owner: Security Owner (interim: Founder/CEO)**

## 1. Data classification

| Class | Examples | Handling |
|---|---|---|
| **Customer Production Data** | scripts, frames, sequences, locks, project metadata | RLS-isolated, encrypted, never used for model training, deletion SLAs below |
| **Account Data** | emails, auth identifiers, billing state | encrypted, minimum retention, Stripe holds payment instruments (never us) |
| **Operational Data** | logs, metrics, cost ledgers | append-only, aggregate reporting only |
| **Public Data** | marketing pages, changelog | no restriction |

## 2. Encryption

- **At rest:** AES-256 on all databases and storage (managed cloud KMS). Customer-managed keys (BYOK) are available on dedicated/VPC deployment tiers, where the platform never holds the unwrap key. Key rotation follows the cloud provider's managed-KMS rotation schedule.
- **In transit:** TLS 1.2+ on every hop — client ↔ platform, platform ↔ database/storage, platform ↔ AI vendors. No plaintext transport paths exist.

## 3. AI-specific data protection

- Customer content is **never used to train AI models** — ours or any vendor's. All upstream inference runs on commercial API terms that exclude training use of inputs and outputs (enforcement mechanics published at aimpreviz.com/security).
- Vendors process content solely to return output. Generated assets persist only in the customer's tenant-scoped storage.
- Per-tenant engine exclusions: any vendor a customer's legal team excludes is blocked at the router; the platform can also pin an engagement to a customer-approved vendor list.

## 4. Retention & deletion

- **Default lifecycle:** customer content persists while the account/project is active. Soft-deleted assets are restorable for 30 days, then purged.
- **On customer deletion request (enterprise engagements):** deletion from production systems begins immediately and completes within **72 hours**; backup copies age out of the point-in-time-recovery window within **30 days**. Written confirmation is provided; cryptographic erasure proof is available for tenants that require it.
- **End of engagement:** customer data is deleted within 30 days of the permitted purpose completing, unless the customer requests export or extended retention in writing.

- **Method:** deletion follows NIST SP 800-88 principles — logical deletion plus provider-level media sanitization (managed cloud storage; we operate no physical media, so no physical destruction process applies).

## 5. Data loss prevention

- Customer production data is confined to the platform: private storage buckets, signed URLs, no public objects, external sharing off by default (and disabled entirely for enterprise engagements that require it).
- Every export, share, and generation event is recorded in the append-only activity log; bulk exports are attributable to an individual account.
- Personnel handling enterprise data do so from devices governed by the Endpoint & Device Security Policy; customer data is not moved to removable media or personal services.

# Patch & Vulnerability Management Policy

**v1.0 · Adopted 2026-07-06 · Owner: Security Owner (interim: Founder/CEO)**

## 1. Surface

The platform is serverless-managed: database, auth, storage, and edge runtime are operated and patched by the managed platform provider (their SOC 2 / infrastructure attestations are inherited and documented). Our patchable surface is application code and its dependency tree.

## 2. Detection

- **Dependency scanning:** automated alerts on vulnerable dependencies (GitHub Dependabot) on the production repository; reviewed weekly.
- **Configuration scanning:** the managed platform's security advisors (RLS/config lints) are reviewed weekly and before every enterprise engagement begins.
- **Penetration testing:** an external vulnerability scan is run before each enterprise engagement's go-live; an annual third-party penetration test is scheduled as part of the SOC 2 program, with independent remediation validation for Critical/High findings.

## 3. Remediation SLAs

| Severity | Fix target |
|---|---|
| Critical (exploitable, customer data at risk) | 48 hours |
| High | 7 days |
| Medium | 30 days |
| Low | 90 days / next scheduled release |

Deployment is continuous (see SDLC policy), so fixes ship the moment they pass review — the SLAs above are ceilings, not cadences.

## 4. Records

Each vulnerability above Medium gets a tracked entry: source, severity, affected component, fix commit, deploy date. Scan reports and pen-test reports are retained for 3 years and shareable under NDA.

# Secure Development (SDLC) Policy

**v1.0 · Adopted 2026-07-06 · Owner: Security Owner (interim: Founder/CEO)**

## 1. Change flow

All production change moves through one pipeline: feature branch → review → `main` → automated deploy (application and edge functions). Database schema changes ship as versioned SQL migrations, reviewed like code and applied deliberately — never ad hoc in production consoles.

## 2. Review requirements

- Every change is reviewed before merge. Changes touching auth, RLS policies, storage paths, billing, or AI-vendor integration receive explicit security review by the Security Owner.
- New tables must ship with RLS policies in the same migration. A table without RLS does not merge.
- Secrets never enter the repository; credentials live in the platform's secret manager. Client code receives only publishable keys.

## 3. Security defaults for new code

- Server-side enforcement first: authorization decisions live in RLS and edge functions, never solely in the client.
- All user-generated URLs pass the URL guard; upstream model calls carry moderation detection; costs are bounded by per-user monthly ceilings.
- Generated assets record provenance (engine, parents) at creation.

## 4. AI-assisted development

AI-assisted code follows the same pipeline — reviewed, typed, built, and tested before merge. The Cortex improvement queue is human-gated: no autonomous proposal reaches production without a human approving the build.

## 5. Release integrity

- Type-checking and production build must pass before merge.
- Deploys are attributable (commit history is the change record) and reversible (previous build redeployable; database restorable via PITR).

# Incident Response Plan

**v1.0 · Adopted 2026-07-06 · Owner: Security Owner (interim: Founder/CEO) Test cadence: tabletop exercise at least annually; first exercise scheduled July 2026 (record kept below).**

## 1. What counts as an incident

Confirmed or suspected: unauthorized access to customer data; credential or key compromise; data exfiltration; malicious code in the pipeline; upstream vendor breach affecting our data; prolonged integrity/availability failure with security cause.

## 2. Roles

- **Incident Commander:** Security Owner (interim: Founder/CEO).
- **Technical lead:** senior engineer on the affected surface.
- **Communications:** Incident Commander owns all customer and external notifications; nobody else communicates externally about an incident.

## 3. Severity

| Sev | Definition | Response start |
|---|---|---|
| SEV-1 | Confirmed customer-data exposure or active compromise | Immediately, 24/7 |
| SEV-2 | Suspected exposure; contained compromise; vendor breach with possible impact | < 4 hours |
| SEV-3 | Attempted attack blocked; vulnerability discovered unexploited | Next business day |

## 4. Response steps

1. **Detect & declare** — via monitoring alerts, security audit log, vendor notice, or report to security@aimpreviz.com. Declare severity; open an incident record (timeline begins).
2. **Contain** — revoke affected credentials/sessions; suspend affected accounts; disable the affected module (per-function isolation allows taking one surface offline while the platform continues); rotate keys.
3. **Assess** — what data, which tenants, what window; preserve logs (append-only activity + security audit logs are the forensic record).
4. **Notify** — see §5.
5. **Eradicate & recover** — fix root cause through the standard pipeline (expedited review), restore from PITR if integrity is affected.
6. **Post-mortem** — written within 5 business days: timeline, root cause, customer impact, corrective actions with owners and dates.

## 5. Notification commitments

- **Affected customers:** notified without undue delay, and **within 24 hours** of confirming or reasonably suspecting an incident involving their data — including what is known, what is being done, and a named

contact.

- **Enterprise engagements with contractual channels** (e.g., a customer security address such as security@amazon.com): notified via that channel within the same 24-hour window, per contract.
- **Regulatory notifications** where applicable follow legal counsel's direction.

## 6. Vendor incidents

Subprocessor register (see Vendor Management Policy) lists each vendor's security-notice channel. A vendor breach touching our data is handled as our incident: same severities, same notification clock.

## 7. History

No data breaches to date (platform history: 2024–present). This section is the permanent record; any future incident's post-mortem is referenced here.

## 8. Exercise log

| Date | Type | Scenario | Findings / actions |
|---|---|---|---|
| 2026-07 (scheduled) | Tabletop | Upstream vendor key compromise | — |

# Security Awareness Training Policy

**v1.0 · Adopted 2026-07-06 · Owner: Security Owner (interim: Founder/CEO)**

- Everyone with production or customer-data access completes security awareness training **within 30 days of access grant** and **annually** thereafter. Completion is recorded (name, date, module).
- Minimum content: phishing and social engineering; credential and MFA hygiene; customer-data handling (classification + deletion policies); incident reporting (what, to whom, how fast); AI-specific rules (no customer data in unapproved AI tools; upstream vendors are contract-bound, personal AI accounts are not).
- Interim delivery (pre-compliance-platform): a recorded internal session + written acknowledgment of this policy pack. Once the compliance platform is live, its training modules replace the interim format and track completion automatically.
- Contractors with any customer-data access are in scope, before access.
- Role supplements: engineers additionally acknowledge the SDLC policy; anyone operating enterprise engagements additionally acknowledges the engagement's data-handling terms.

# Vendor Management Policy & Subprocessor Register

**v1.0 · Adopted 2026-07-06 · Owner: Security Owner (interim: Founder/CEO)**

## 1. Policy

- Every vendor that stores, processes, or transmits customer data is listed in the register below before first production use.
- **AI vendors:** commercial API tiers only, with terms excluding training use of inputs and outputs. Term snapshots are collected and retained at onboarding and re-checked at renewal or on public terms changes.
- Vendors are reviewed **annually**: security posture (attestations where published), terms changes, incident history, continued need.
- Customer-facing controls: per-tenant engine exclusions are enforced at the router; an enterprise engagement can be pinned to a customer-approved vendor list. Enterprise customers are notified of material subprocessor additions affecting their engagement.
- Exit: each vendor is swappable behind the orchestration layer; no single AI vendor is load-bearing for the platform's availability.

## 2. Subprocessor register

| Vendor | Function | Customer data reaches it? | Data type | Training use |
|---|---|---|---|---|
| Supabase (on AWS) | Database, auth, storage, edge functions | Yes — primary store | All platform data | N/A (infrastructure) |
| xAI | Image generation/editing; agent reasoning | Yes — inference only | Prompts, frames | Contractually excluded |
| Google (via AIM AI Gateway) | Image/vision models; drift scoring; planning | Yes — inference only | Prompts, frames | Contractually excluded |
| fal.ai | Image/video/3D inference (multi-model) | Yes — inference only | Prompts, frames | Contractually excluded |
| ElevenLabs | Voice synthesis | Yes — inference only | Voice text/refs | Contractually excluded |
| Stripe | Payments | Payment data only (never touches our DB) | Card/billing | N/A |
| Resend | Transactional email | Email addresses | Account email | N/A |
| Vimeo | Marketing video hosting | No customer production data | Public assets | N/A |

Every generated asset records which engine produced it (Lineage), so the subprocessors that touched any given frame are attributable per asset — auditable evidence, not policy assertion.

## 3. AI provider detail (for GenAI security reviews)

For each AI vendor above: data leaves our system boundary solely for inference; the vendor returns the output and does not retain content for training; short-lived abuse-prevention logs, where a vendor keeps them, carry the same no-training terms. Enforcement mechanics are published at aimpreviz.com/security (Zero Training Retention section).

# Business Continuity & Disaster Recovery Plan

**v1.0 · Adopted 2026-07-06 · Owner: Security Owner (interim: Founder/CEO)**

## 1. Architecture posture

The platform runs on managed, multi-AZ cloud infrastructure (Supabase on AWS): the database, storage, auth, and edge runtime are operated by the platform provider with their own attested continuity controls, which we inherit and document. Application code and infrastructure-as-code (migrations) live in version control and are redeployable to a fresh project from the repository.

## 2. Backups

- **Database:** automated daily backups plus point-in-time recovery (WAL), giving a recovery point of minutes for the recent window and ≥ 30 days of restore depth. Managed and encrypted (AES-256) by the platform.
- **Storage (assets):** durable object storage with provider-level redundancy; asset provenance (Lineage) allows regeneration paths for derived assets.
- **Backup failure alerting:** platform backup status is checked weekly and on-call alerts are raised through the system-alert channel; failures are treated as SEV-2 incidents.

## 3. Objectives

| Metric | Target |
|---|---|
| RPO (data loss ceiling) | ≤ 15 minutes (PITR window) |
| RTO — database restore | ≤ 4 hours |
| RTO — full platform on fresh infrastructure | ≤ 24 hours (repo + migrations + secrets restore) |

## 4. Scenarios & playbooks

1. **Data corruption / bad migration:** PITR restore to pre-event timestamp; replay reviewed migrations; verify with integrity checks.
2. **Region/provider outage:** monitor provider status; if extended, restore latest backup to alternate region/project; repoint DNS; rotate secrets.
3. **Account/credential compromise:** see Incident Response Plan (takes precedence); recovery via restore after containment.
4. **Key-person unavailability:** credentials escrowed so that two officers can jointly recover admin access; runbooks in this pack are written to be executable by any senior engineer.

## 5. Testing

- A **restore test** (database PITR restore to a scratch project, verified row counts + spot integrity checks) is performed at least **annually** and before any enterprise engagement go-live; results are recorded below.
- The first recorded test is scheduled for July 2026 as part of the current enterprise-engagement preparation.

## 6. Test log

| Date | Test | Result | RTO/RPO achieved | Issues |
|---|---|---|---|---|
| 2026-07 (scheduled) | DB PITR restore to scratch project | — | — | — |

# Logging & Monitoring Policy

**v1.0 · Adopted 2026-07-06 · Owner: Security Owner (interim: Founder/CEO)**

## 1. What is logged

- **User activity log (append-only):** generation, edit, export, share, and billing events — attributable to an individual account, timestamped.
- **Security audit log (append-only):** authentication-relevant and administrative actions (role changes, credit adjustments, suspensions, refunds), with actor and target.
- **Per-generation records:** every AI job records engine, parameters, latency, outcome, and cost; every asset records provenance (Lineage).
- **Routing decisions:** engine-selection decisions for locked renders are ledgered (router_decisions).
- **Platform layer (inherited):** auth events, API logs, and database logs are captured by the managed platform.

Logs never contain plaintext passwords, tokens, or payment data.

## 2. Retention

- Application security/audit/activity logs: retained **≥ 13 months**; append-only by design (no user-facing update/delete paths; writes are service-role only).
- ≥ 90 days immediately queryable (hot) — in practice the full history is SQL-queryable at all times.
- Logs live in the managed database tier, operationally separate from the application runtime that produces them, and are covered by the backup regime (daily + PITR).
- Platform-layer logs follow the provider's plan-tier retention and are exportable during investigations.

## 3. Monitoring & alerting

- System alert checks run scheduled health evaluations (generation failure rates, queue depth, cost anomalies) and raise alerts to the operations channel.
- Monthly cost ceilings bound abuse blast-radius per account.
- Security-relevant alerts (auth anomalies, backup failures, vendor notices) are triaged under the Incident Response Plan's severity matrix.

## 4. Customer access

For enterprise engagements, relevant security-event extracts are provided to the customer on request during security investigations, within 5 business days (faster during an active incident, per the IR plan).

# Endpoint & Device Security Policy

**v1.0 · Adopted 2026-07-06 · Owner: Security Owner (interim: Founder/CEO)**

Applies to every device used to access production systems or enterprise customer data (including contractor devices — see BYOD below).

## 1. Device standards

- Full-disk encryption enabled (FileVault / BitLocker).
- OS auto-updates on; current major OS version or n-1.
- Endpoint protection active with automatic updates (platform-native — XProtect/ Defender — minimum; EDR tooling arrives with the compliance platform).
- Screen lock required, auto-lock at ≤ **15 minutes** idle.
- Daily work under a non-administrator account; local admin reserved for installs.
- Device registered with the Security Owner (owner, model, serial) before production access.

## 2. Data handling on devices

- Enterprise customer data is worked in the platform, not on laptops: no copies to removable media, personal cloud storage, or personal messaging.
- Browser sessions to production consoles are MFA-protected; credentials live in a password manager, never in plaintext files.

## 3. BYOD

Permitted only under this same standard, registered the same way, and subject to remote credential revocation (session + password reset) on loss or departure. Devices that can't meet the standard don't touch enterprise data.

## 4. Loss or compromise

Report immediately (security@aimpreviz.com); treated per the Incident Response Plan — sessions revoked and credentials rotated first, questions after.

## 5. Attestation

Interim: each person signs an annual attestation of compliance with this policy; the compliance platform's device-management agent will replace self-attestation with continuous verification when onboarded.

# System Architecture & Data Flow

**v1.0 · 2026-07-06 · Owner: Security Owner (interim: Founder/CEO) Purpose:** the comprehensive architecture and data-flow description for enterprise security review (systems, boundaries, GenAI entry/exit points, storage locations).

## 1. Components

| Component | Role | Customer data |
|---|---|---|
| Browser client (React SPA) | UI only; holds a short-lived JWT and publishable keys | In session memory only |
| Edge functions (Deno, server-side) | All business logic: auth checks, billing, AI orchestration, logging | Processes; does not persist locally |
| Postgres (managed, AWS) | System of record; RLS on every table | Stored, AES-256 at rest |
| Object storage (managed, AWS) | Generated assets, uploads; tenant-/project-scoped paths, private buckets, signed URLs | Stored, AES-256 at rest |
| Auth layer (GoTrue) | Accounts, MFA (TOTP), sessions; bcrypt password hashing | Identifiers only |
| AI vendors (see subprocessor register) | Inference only, via server-side keys | Transits for inference; not retained for training |
| Stripe | Billing | Payment data never enters our systems |

## 2. Data flow (happy path)

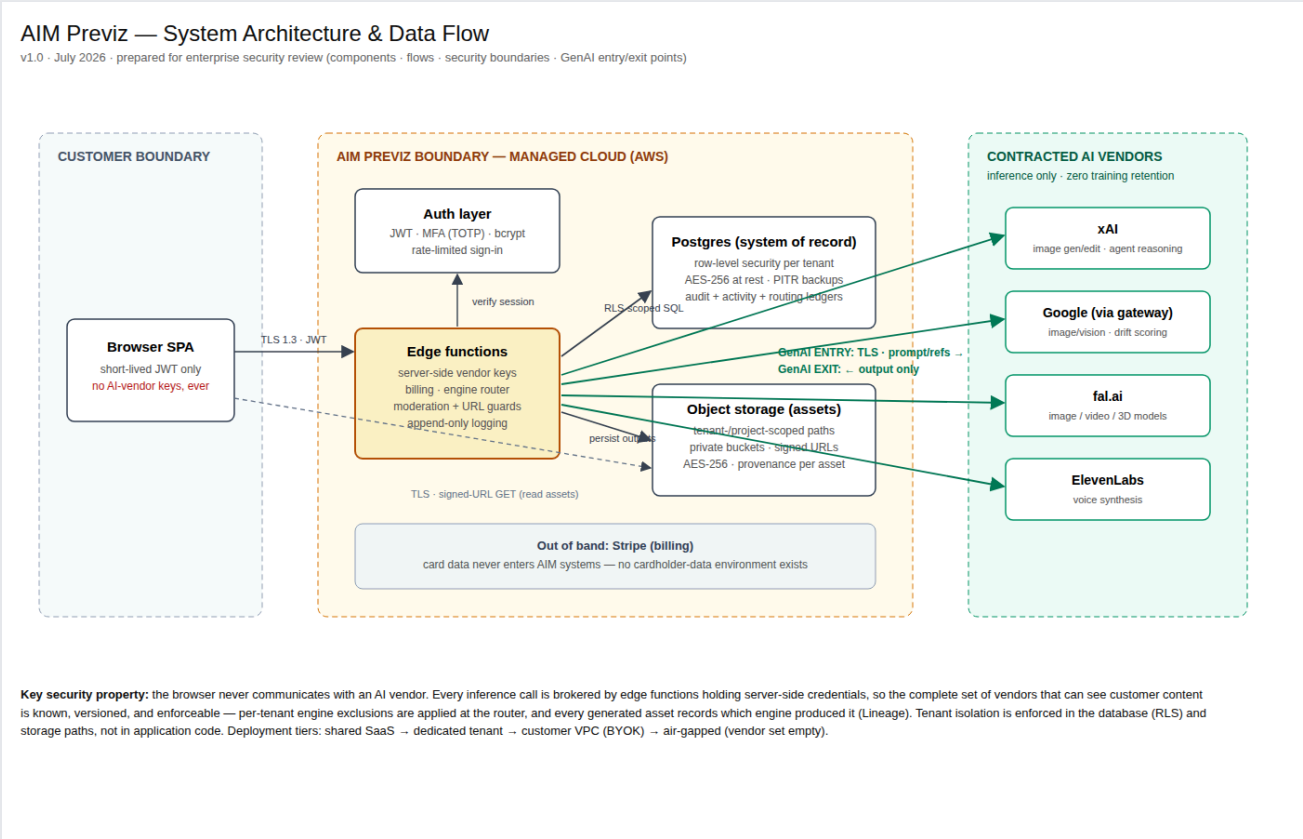


Rendered system diagram (v1.0). Editable source lives in the repository (docs/security/architecture-data-flow.md, Mermaid).

**Key property: the browser never talks to an AI vendor.** Every inference call is brokered by edge functions holding server-side credentials, so the complete set of vendors that can ever see customer content is known, versioned, and enforceable (per-tenant exclusions at the router).

## 3. GenAI entry & exit points (per enterprise GenAI review requirements)

- **Entry:** a prompt/reference leaves the AIM boundary only at the edge function → vendor API hop (TLS), only to vendors on the subprocessor register, only for the requesting tenant's job.
- **Exit:** the vendor returns the generated output to the edge function, which persists it into the requesting tenant's scoped storage and records provenance (engine + parents) on the asset. Continuity drift-scoring runs the same path (gateway vision model), server-side.
- **No secondary paths:** no vendor webhooks into customer data, no client-side vendor SDKs, no training/fine-tuning pipelines — ours or vendors' — touch customer content.

## 4. Security boundaries

1. **Identity boundary:** JWT + MFA at the auth layer; every request carries an individual identity.
2. **Tenant boundary:** enforced in the database (RLS) and storage paths — not in application code. Backups inherit the same structure.
3. **Vendor boundary:** server-side keys; inference-only contracts; router exclusions; per-asset engine attribution (Lineage).

4. **Billing boundary:** Stripe-hosted; card data never enters AIM systems — no cardholder-data environment exists in this architecture.

## 5. Deployment tiers

Shared SaaS (above) · Dedicated tenant (separate DB/storage/functions) · Customer VPC (customer-managed keys, egress allowlist) · Air-gapped (open-weight models on customer hardware; vendor set is empty).

## 6. Detection layers

Append-only activity + security audit logs, per-generation records, routing ledger, scheduled system-alert health checks, monthly cost ceilings, backup monitoring — see Logging & Monitoring Policy.

# Evidence — Encryption in Transit (TLS)

**Captured: 2026-07-06 · Method: live TLS handshake inspection (curl -v) · Serves: Q8**

## Application endpoint — aimpreviz.com

```
SSL connection using TLSv1.3 / TLS_AES_256_GCM_SHA384 / X25519 / id-ecPublicKey
subject: CN=aimpreviz.com
expire date: Sep  2 2026 GMT
```

## Data-plane endpoint (managed database/API gateway)

```
SSL connection using TLSv1.3 / TLS_AES_256_GCM_SHA384 / X25519 / id-ecPublicKey
```

## Interpretation

- Both the application origin and the data plane negotiate **TLS 1.3** with **AES-256-GCM** — exceeding the TLS 1.2 minimum requirement.
- HTTP endpoints redirect to HTTPS; HSTS is served by the edge (Cloudflare).
- Server-to-vendor inference calls use HTTPS exclusively (see architecture-data-flow: no plaintext transport paths exist).

Re-verification: this handshake can be reproduced at any time with `curl -svI https://aimpreviz.com` or an SSL Labs scan of aimpreviz.com.

# Evidence — Tenant Isolation via Row-Level Security

**Captured: 2026-07-06 · Method: static analysis of the versioned migration history · Serves: Q7, Q11, Q17**

## Quantified posture

- **361** `CREATE POLICY` statements across the migration history
- **84** tables with `ENABLE ROW LEVEL SECURITY`
- Policy-per-table is a merge requirement: per the Secure Development Policy, a new table does not merge without RLS in the same migration.

## Representative policies (verbatim from migrations)

```
CREATE POLICY "Users can view their own projects"
  ON public.projects FOR SELECT USING (auth.uid() = user_id);

CREATE POLICY "Users can create their own projects"
  ON public.projects FOR INSERT WITH CHECK (auth.uid() = user_id);

CREATE POLICY "Users can update their own projects"
  ON public.projects FOR UPDATE USING (auth.uid() = user_id);
```

## Deny-all pattern for sensitive tables

The security audit log is intentionally **stricter than select-own**: RLS is enabled with **no user policies at all**, so only the service role (server-side edge functions) can touch it:

```
ALTER TABLE public.security_audit_log ENABLE ROW LEVEL SECURITY;
-- No policies = deny all for anon/authenticated
-- Audit logs are accessed only via admin edge functions
```

## Storage isolation

Asset objects live under tenant-/project-scoped paths (`{user_id}/{project_id}/…`) in private buckets; access is via signed URLs only. Backups inherit the same structure.

# Evidence — Security Event Logging Schemas

**Captured: 2026-07-06 · Method: verbatim schema extracts from versioned migrations · Serves: Q9, Q10**

## 1. User activity log (append-only)

```
CREATE TABLE public.user_activity_log (
  id uuid PRIMARY KEY DEFAULT gen_random_uuid(),
  user_id uuid NOT NULL,
  action_category text NOT NULL,  -- 'generation', 'export', 'share', 'billing', ...
  action_type text NOT NULL,
  project_id uuid REFERENCES public.projects(id) ON DELETE SET NULL,
  details jsonb DEFAULT '{}',
  ai_model text,
  ai_provider text,
  credits_consumed numeric,
  success boolean DEFAULT true,
  error_message text,
  -- immutable audit record; writes via service role only
  ...
);
```

## 2. Security audit log (append-only, service-role only)

```
CREATE TABLE public.security_audit_log (
  id uuid PRIMARY KEY DEFAULT gen_random_uuid(),
  user_id uuid REFERENCES auth.users(id) ON DELETE SET NULL,
  action text NOT NULL,
  resource_type text NOT NULL,
  resource_id text,
  details jsonb DEFAULT '{}'::jsonb,
  ip_address text,
  user_agent text,
  success boolean NOT NULL DEFAULT true,
  created_at timestamptz NOT NULL DEFAULT now()
);
-- RLS enabled with no user policies: deny-all except service role
```

## 3. Additional ledgers

- **Per-generation job records** — engine, parameters, latency, outcome, cost (infinite_jobs + cost logger).
- **Engine routing ledger** — router_decisions: default vs chosen engine, deciding basis, failover outcome, per decision.
- **Continuity drift records** — continuity_checks: per-render quality grading including the producing engine.

### Retention properties

- No user-facing UPDATE/DELETE paths exist on log tables (append-only by schema + RLS design).

- All history is SQL-queryable (exceeds the 90-day hot requirement); policy retention floor is 13 months (Logging & Monitoring Policy §2).
- Log tables live in the managed database tier — separate from the edge runtime that produces events — and are covered by daily backups + PITR.

# Evidence — Authentication Flow & Credential Handling

**Captured: 2026-07-06 · Method: production login flow screenshot + auth-layer documentation · Serves: Q1, Q2, Q12**

## Login flow (production build)

Screenshot: `login-flow.png` (this pack). The flow is email + password with "Forgot password" reset; Google/Apple OAuth are visible as **coming soon** (not yet enabled — stated accurately). A consent banner covers GDPR/CCPA/LGPD cookie categories.

## Credential handling

- Passwords are stored only as **one-way bcrypt hashes** by the auth layer (GoTrue); plaintext is never stored or logged.
- Failed sign-ins are rate-limited at the auth layer.
- Sessions are short-lived JWTs with automatic expiry/refresh and server-side revocation; no long-lived static tokens.
- MFA (TOTP) is available at the auth layer and is **enforced for enterprise engagement operator accounts** per the Access Control Policy §3.
- Enterprise SSO (SAML 2.0 / OIDC) is available on the dedicated-tenant deployment tier.

## Internal systems

MFA is required on all consoles handling customer data or production configuration (Supabase, GitHub, Google Workspace, Stripe) — evidence screenshots of each console's enforcement setting are collected in the engagement evidence folder at submission time.

# Evidence — Data Deletion Mechanics

**Captured: 2026-07-06 · Method: shipped product behavior + code references · Serves: Q15, Q4**

## Shipped deletion behavior

- **Soft-delete with bounded restore:** deleted canvas assets move to a Trash bin, restorable for 30 days, then **purged automatically** — shipped product behavior (public changelog, "Previz Infinite — a 30-day Trash bin for deleted assets").
- Agent-initiated discards follow the same path by rule: soft-delete to trash, never silent hard deletion (code: useInfiniteAgent discard verb).
- **Hard deletion path (enterprise/on-request):** account/project deletion removes database rows and storage objects; backup copies age out of the point-in-time-recovery window within 30 days.

## SLAs (Data Protection Policy §4)

- Customer deletion request: production deletion completes within **72 hours**; written confirmation provided; cryptographic erasure proof available.
- End of engagement: deletion within 30 days of purpose completion.
- Method follows NIST SP 800-88 (logical sanitization + provider-level media controls; no physical media operated).

# Evidence — AI Vendor Terms (No-Training) Reference Links

**Captured: 2026-07-06 · Serves: Q5(A)(E) · Companion to the Subprocessor Register**

Canonical public terms pages for each inference vendor. Term snapshots (PDF) are collected at vendor onboarding and re-checked at renewal; the links below are the vendors' authoritative current versions.

| Vendor | Function in AIM Previz | Terms / data-use page |
|---|---|---|
| xAI | Image generation/editing; agent reasoning | https://x.ai/legal |
| Google (Gemini API, via gateway) | Image/vision; drift scoring; planning | https://ai.google.dev/gemini-api/terms |
| fal.ai | Image/video/3D inference | https://fal.ai/terms |
| ElevenLabs | Voice synthesis | https://elevenlabs.io/terms-of-use |
| Supabase | Infrastructure (DB/auth/storage/functions) | https://supabase.com/security |
| Stripe | Billing (card data never enters AIM systems) | https://stripe.com/privacy |

All inference runs on commercial API tiers whose terms exclude training use of inputs and outputs; enforcement mechanics are published at https://aimpreviz.com/security (Zero Training Retention section). Per-asset engine attribution (Lineage) makes vendor exposure auditable per frame.