Protocol Wealth — Privacy, Security & Compliance Posture

v0.1 — 2026-05-19 snapshot

Effective: 2026-05-19
Owner: Nick Rygiel (CTO/CISO) + Adam Blumberg (CCO)

Contents

Protocol Wealth — Privacy, Security & Compliance Posture

(2026-05-19 snapshot)

Protocol Wealth (SEC RIA, CRD #335298) operates an AI-native investment advisory practice. This document summarizes the privacy, security, and compliance substrate we have built — including the engineering layers most firms do not surface to public view.

At a glance

Layer	What it does	Where it lives
Zero Data Retention (ZDR) workspace	Anthropic ZDR — contractual and enforced by the vendor at the API workspace level; our content is not retained for training	All AI-using surfaces
Schema-level PII tagging (`pii.{high,medium,low}`)	Fields tagged at ingestion; high-tag fields are structurally excluded from any LLM-bound payload	pw-api + downstream BFFs
Independent PII egress canary	Second-layer guard at every Anthropic SDK call; pattern set deliberately re-implemented byte-identical across surfaces	pw-os-v2 + pw-portal-v2 (pw-api wave queued)
WORM audit retention	7-year retention-locked Google Cloud Storage bucket mirrors the application audit log	Per SEC Rule 17a-4(b)(4) + 17a-4(f)(2)(ii)
Sentinel-row reconciliation	Reconciliation against immutable evidence rows emits new linked rows; never UPDATE	audit_log, KYC, e-signature archive
Canonical webhook receiver	Every vendor callback flows the same six stages: verify, dedup, parse, process, audit, dead-letter	All vendor integrations

Layer	What it does	Where it lives
Defense-in-depth onboarding	Veriff identity, Scorechain AML, Anvil e-signature — each via reviewed webhook handler	New-client onboarding
Single-cloud sovereignty	Google Cloud only (ISO 27001 / SOC 2 alignment); no multi-cloud client-data spread	All PW-internal services

---

System architecture (abridged)

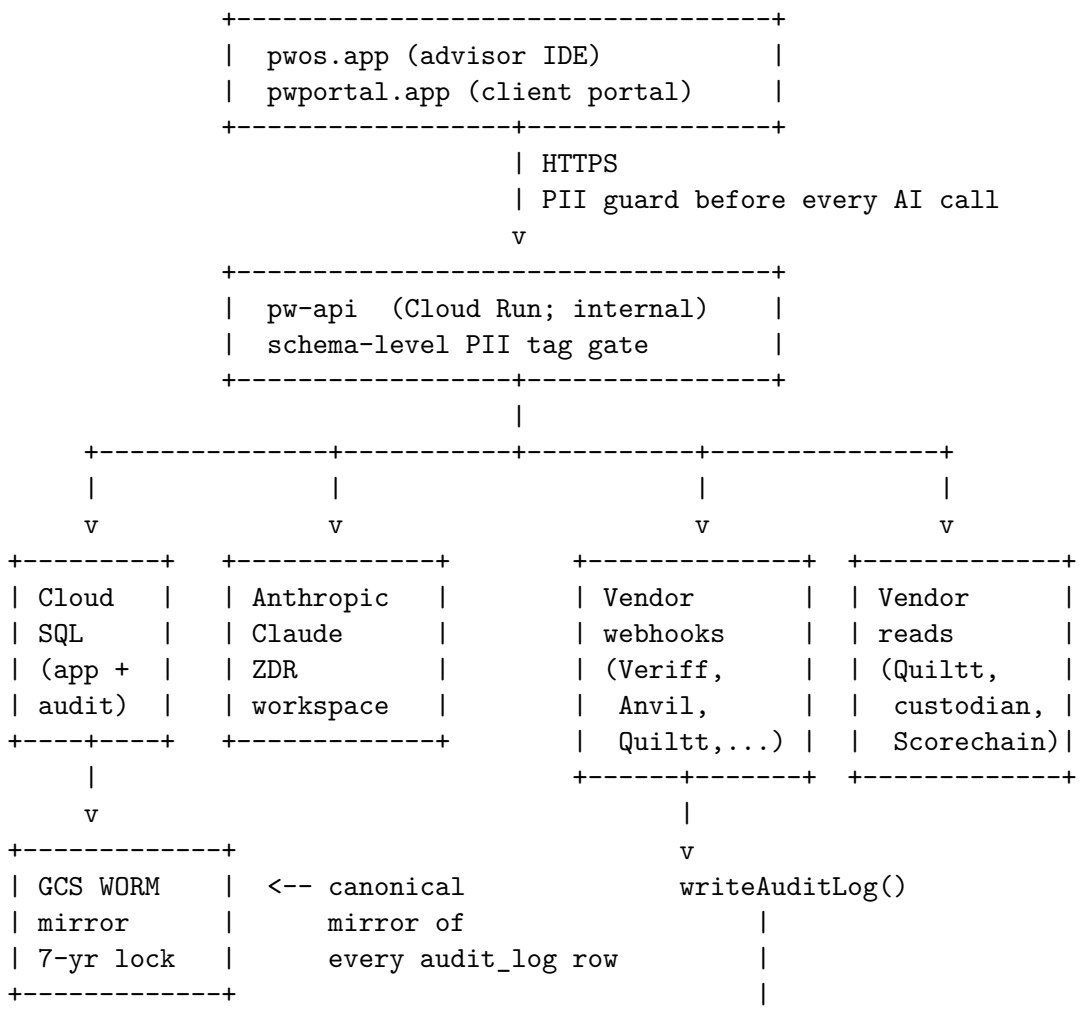

All client data ingestion, AI inference, and audit retention occurs on infrastructure located within the United States.

Privacy

Reg S-P §248.30(b) safeguards plus a layered AI-privacy substrate.

Zero Data Retention with Anthropic

We process AI workloads exclusively through Anthropic's Claude API under a Zero Data Retention agreement formally approved on April 21, 2026. The ZDR posture is enforced by Anthropic at the API workspace level (workspace ID `fe30b317-0a9c-47df-a972-b9611e6e0002`) — not by our internal policy alone. Inputs and outputs are not retained beyond the duration needed to process a request, inference runs on US-based infrastructure, and our content is contractually excluded from training, fine-tuning, or model improvement.

Three-bucket PII tagging at the schema level

Every client-data field in our application database carries a `pii.high`, `pii.medium`, or `pii.low` tag applied at ingestion. The taxonomy is codified in an Architecture Decision Record and enforced by middleware:

- `pii.high` — government identifiers (SSN, DOB), account-control fields (full account numbers, wallet seed phrases, private keys), biometric data, authentication artifacts. **Fail-closed:** stripped from every LLM-bound payload unless a single-use field-level waiver token explicitly authorizes a specific field path.
- `pii.medium` — quasi-identifiers and financial context (email, full street address, employer name, balance ranges, transaction memos). Pass-through with audit logging.
- `pii.low` — non-identifying public context (first name, state of residence, age bracket, generic occupation, broad asset-class allocation). Pass-through; not logged at field level.

The taxonomy lives in a single canonical map at `pw-api/src/lib/pii-tags.ts`; downstream services hold byte-equal derived copies, and a CI drift-check enforces alignment.

Independent PII egress canary

A second-layer guard runs at every Anthropic SDK call site. The canary's pattern set is deliberately re-implemented from scratch in byte-identical copies across surfaces — pw-os-v2 and pw-portal-v2 today, pw-api on the next iteration — so the two layers cannot share a bug. Any residual PII detected at egress raises `PIIEgressCanaryError`, aborts the call, and surfaces as a structured

`pii_blocked` event. Cloud Logging metrics and alerts fire on every blocked event and on any detector-mismatch (canary scanned fewer turns than the request carried).

Data never sold; never used to train

We do not sell client information. Client data is not used to train AI models. Third-party service providers (custodians, data aggregators, CRM) are contractually restricted to using your data only to provide services to us.

Encryption + access discipline

- **In transit:** TLS 1.2 minimum (TLS 1.3 preferred) externally; TLS 1.3 internally.
- **At rest:** AES-256 with Google-managed keys; sensitive fields additionally encrypted at the application layer.
- **Secrets:** Google Secret Manager with IAM-scoped access and full audit logging; no long-lived credentials in source.
- **Multi-factor authentication** required for all advisor and client surfaces; passkey-first authentication on new onboarding flows.
- **Least-privilege IAM** per service; workload identity federation for CI; no static service-account keys.

Security substrate

Information Security Program (WISP)

A written information security program covers administrative, technical, and physical safeguards under SEC Regulation S-P §248.30(b). The WISP is reviewed quarterly and is current as of v1.2.

Incident Response Plan

An incident response plan (v1.2) covers detection, containment, eradication, recovery, and post-incident review. Breach notification readiness anchors to the 30-day notification window under amended Reg S-P. Tabletop exercises are scheduled quarterly starting Q2 2026.

Infrastructure controls

- **Cloud Run** services run in isolated containers with automatic security patching and pinned revisions.
- **Cloud SQL (Postgres)** is private-network-only with IAM authentication, automatic encryption, automated backups, and row-level security enforced by the database engine itself (isolation is structural, not application-convention).
- **Memorystore Redis** is private-network-only.
- **Cloud Audit Logs** are immutable.

- **Cloudflare** fronts the public marketing site only; authenticated advisory surfaces (pwos.app, pwportal.app, pw-api) route direct-to-GCP without edge interference.

Monitoring + alerting

Cloud Logging metrics fire on PII egress canary blocks, detector-mismatches, ZDR configuration drift, audit-mirror write failures, and webhook signature failures. Alerts route to an operator channel with documented runbooks. The full chain is end-to-end: application canary aborts on residual PII → structured log → metric → alert → runbook → CCO loop-in.

Boundary discipline

PW operates entirely on Google Cloud Platform — a deliberate single-cloud posture for ISO 27001 / SOC 2 alignment and data sovereignty. Cross-provider dependencies are introduced only where structurally required (e.g., qualified custodians) and never for client-data substrate.

Compliance posture

Regulatory anchors

- **SEC-registered investment adviser**, CRD #335298. Form ADV current.
- **Investment Advisers Act of 1940, Rule 204-2** (Books and Records) — advisory records retained at least 5 years, with the most recent 2 years in easily accessible locations. We retain audit-log records for 7 years as a conservative default.
- **SEC Regulation S-P** (Privacy of Consumer Financial Information) — written program; safeguards rule per §248.30(b); 30-day breach notification commitment per the amended rule.
- **SEC Regulation S-P / Reg XP** small-entity tier — June 3, 2026 readiness on track (IRP, WISP, Subprocessors, Security Posture, KYC/AML, Customer Breach Notification Template, Reg S-P Compliance Record 2026).
- **SEC Rule 17a-4** (Records Preservation) — WORM mirror substrate with bucket-level retention lock satisfies 17a-4(f)(2)(ii) electronic-storage requirements; STANDARD → NEARLINE lifecycle aligned to the 2-year immediate-access horizon of 17a-4(f)(3).
- **Marketing Rule, §206(4)-1** — every client-facing communication routes through CCO sign-off; AI-assisted advisor drafts are never delivered to clients directly.
- **Gramm-Leach-Bliley Act (GLBA)** — annual privacy notice in standard FACTS format.
- **State privacy laws** — California (CCPA/CPRA), Colorado, Connecticut, Utah, Virginia, and other comprehensive consumer privacy regimes honored.
- **ESIGN Act + UETA** — electronic signature flow (Anvil) produces tamper-evident PDF/A records with completion certificates retained alongside the source documents.

AI governance

Every AI-assisted output is attributable to a named foundation model at Anthropic. Prompts, model identifiers, response content, and timestamps are logged subject to the PII redaction controls described above. The single-vendor transparency posture is deliberate: a regulator or qualified partner evaluates one upstream AI relationship rather than a chain of undisclosed subprocessors.

Three model aliases (`CLAUDE_MODEL_FRONTIER`, `CLAUDE_MODEL_WORKHORSE`, `CLAUDE_MODEL_LIGHTWEIGHT`) abstract model selection from call sites; an ESLint rule blocks hardcoded model strings firmwide.

Co-intelligence framework

Human advisers and AI work together; neither operates alone on matters material to a client account. AI assists with research, monitoring, drift detection, and document preparation. Final investment decisions, fiduciary judgment, and all client-facing communications remain with the human adviser.

Audit + recordkeeping discipline

Canonical audit log

Every state-changing operation across the platform emits a row in a canonical `audit_log` table. Action verbs follow a structured three-segment pattern (`<domain>.<entity>.<verb>` — e.g., `pii.field.excluded`, `anvil.envelope.signed`, `veriff.identity_verification.completed`) enforced by a regex. Each row carries actor identity, principal chain (advisor → AI session → tool), resource ID, masked detail, IP address, request ID, trace ID, and a 7-year retention horizon.

WORM mirror

Every `audit_log` row mirrors to a write-once-read-many GCS bucket (`gs://pwllc-audit-archive`) with bucket-level retention lock at 7 years. The lock is irreversible for the retention window — no PW operator, including those with full GCP organization-admin, can delete or modify an object within the window. Each object carries a SHA-256 content hash independently re-computable by a downstream auditor; the source git revision is captured in object metadata.

Sentinel-row reconciliation

The `audit_log` table is immutable at the database trigger level (the BEFORE-UPDATE trigger raises). Reconciliation of failed mirror writes therefore cannot UPDATE the failed row; instead, a daily cron emits a NEW sentinel row (`audit.mirror.retry.success` or `audit.mirror.retry.failure`) referencing the failed row's ID. A recursion guard prevents sentinel rows from triggering their own sentinel chain. The pattern is architectural — applicable to every immutable evidence table we operate, including KYC verifications and the signed-document archive.

Cross-component canonical actions

State changes that span multiple components (e.g., KYC verification completing both a Veriff webhook landing and a downstream risk-tolerance gate progression) emit canonical actions in both namespaces, joined by a shared correlation ID. Cross-component tracking is mechanical, not per-route judgment.

Engineering + operational rigor

Architecture Decision Records

Material substrate decisions are codified as ADRs under `shared/architecture/decisions/`. Each ADR carries status, context, decision, consequences, alternatives considered, and explicit gates (CTO and CCO where compliance-adjacent). Current ADRs cover PII tagging, the WORM audit mirror, the webhook receiver primitive, BFF authentication, the KYC state machine, AML decisioning, the signed-document state machine, the risk-tolerance state machine, and the review-items primitive.

Structural-over-disciplinary

Where a control can be enforced by code, it is enforced by code. The PII egress guard does not rely on every developer remembering to scrub a payload; middleware enforces it and CI lints any LLM SDK call routed outside the middleware. The canonical action catalog is regex-enforced; the model alias map is ESLint-enforced; the PII tag map is drift-check-enforced across repos. The credential-layer thesis (PII never sent to any LLM) is enforced by code that runs regardless of caller behavior, not by caller behavior itself.

Test discipline

Every regulated surface carries integration tests at the route level and unit tests at the substrate level. The egress canary fields are pinned by unit tests across repos so independent canary copies cannot silently drift. Webhook receivers have mock-vendor signature, dedup, parse, process, and dead-letter coverage.

Canary deploys + fail-closed posture

ZDR configuration is verified at boot via a startup assertion that exits the process if the workspace environment variable is missing. Cloud Run holds traffic on the prior healthy revision when an assertion fails — a misconfiguration cannot reach client traffic. Quarterly ZDR configuration audits verify the workspace continues to reflect zero-day retention, US-only inference, and no-training posture.

Multi-agent AI orchestration

Agent-based development runs under a documented coordination protocol (`shared/CLAUDE.md` Agent Coordination Protocol). Hard-stops gate agent authority on terraform, IAM, secrets, migrations, production network, and any compliance interpretation; agents propose and humans dispose on regulated surfaces. A three-strikes rule prevents autonomous retry loops on failing code paths.

Vendor due diligence

Every subprocessor passes through a documented due-diligence framework aligned to Reg S-P §248.30 vendor management expectations. We collect on-file SOC 2 reports, DPA executions, and (where applicable) penetration-test reports; vendor risk assessments refresh annually and on material change. Per-vendor architecture notes live under `shared/architecture/api/<vendor>.md`.

The current active vendor roster (full disclosure: `protocolwealthllc.com/subprocessors`):

Vendor	Role	Attestations on file
Anthropic, PBC	AI inference under ZDR	SOC 2 Type II, ISO 27001
Google Cloud Platform	Compute, storage, secrets, audit logs	SOC 1/2/3, ISO 27001/27017/27018/27701, PCI DSS, FedRAMP High
Cloudflare	DNS, CDN, WAF for public surfaces	SOC 2 Type II, ISO 27001
Veriff OU	Identity verification (KYC)	SOC 2 Type II, ISO 27001
Scorechain S.A.S. (via QuickNode)	OFAC sanctions screening + KYT risk scoring	Under active vendor-risk review
QuickNode, Inc.	Multi-chain RPC + Scorechain integration substrate	SOC 2 Type II
Hadrius, Inc.	AI compliance monitoring + supervision	Under active vendor-risk review
Quiltt, Inc. (with MX, FinGoal)	Financial account aggregation	SOC 2 Type II
Altruist Financial LLC	Advisory custodian + billing	SEC/FINRA oversight, SOC 2 Type II
Interactive Brokers LLC	Brokerage + custody	SEC/FINRA registered broker-dealer
Anchorage Digital Bank, NA	Qualified digital asset custodian	OCC oversight, SOC 2 Type II
BitGo Trust Company	Qualified digital asset custodian	SD Banking oversight, SOC 2 Type II

Vendor	Role	Attestations on file
Fordefi	MPC wallet infrastructure	SOC 2 Type II
Anvil	E-signature with ESIGN/UETA attestation + PDF/A archival	Vendor DD on file
Postmark	Transactional email	SOC 2 Type II
Wealthbox	CRM	SOC 2 Type II

All vendor relationships carry contractual restrictions on data use, breach-notification clauses (72-hour closure where applicable), and US-region processing commitments. Authenticated advisory surfaces do not transit Cloudflare; the public-edge layer handles marketing properties only.

What is open source

We publish a portion of our substrate work as open source under `pwos-core` and adjacent repositories — canonical patterns, ADRs, dispatch infrastructure, and select reusable agent skills. The shipping cadence is PR-by-PR on the underlying shared repository; licensing follows the project README. Contribution guidelines are documented at the project root.

The intent is not marketing surface — it is to demonstrate, in code, the structural-over-disciplinary patterns this document describes.

How to verify any of this

We encourage qualified prospects, partners, and reviewers to verify our posture rather than accept it at face value:

- **Confirm ZDR with Anthropic.** Our workspace ZDR status (effective April 21, 2026; workspace `fe30b317-0a9c-47df-a972-b9611e6e0002`) is attestable by Anthropic on request to an appropriately credentialed reviewer under NDA.
- **Review the substrate.** Architecture Decision Records covering PII tagging, the WORM audit mirror, and the webhook receiver primitive are available for review by qualified institutional prospects under NDA. Public high-level summaries live in this document.
- **Sample audit log + WORM mirror.** Representative anonymized samples of audit-log rows demonstrating action verb structure, principal chain capture, and WORM-mirror round-trip are available for institutional due-diligence reviewers.
- **Vendor DD bundle.** SOC 2 reports, DPAs, and vendor risk assessments are available under NDA for qualified institutional prospects.

- **Written attestation.** A formal written attestation of current AI posture (subprocessor list, ZDR status, inference region, PII controls in effect at time of request) is available on letterhead from the CCO on qualified request.
- **Form ADV Part 2A.** Public at `adviserinfo.sec.gov/firm/brochure/335298`.

Contact pathways

- **Privacy + data handling:** Adam Blumberg, CCO — compliance@protocolwealthllc.com
- **Security + substrate engineering:** Nick Rygiel, CTO/CISO — nick@protocolwealthllc.com
- **Coordinated vulnerability disclosure:** security@protocolwealthllc.com
- **General + scheduling:** https://protocolwealthllc.com (Calendly link in footer)

What we commit to

If something goes wrong:

- We will detect and respond quickly. We have continuous monitoring, automated alerting on the egress canary and audit-mirror substrate, and documented incident response procedures.
- We will notify affected clients within 30 days if sensitive information is compromised, consistent with amended SEC Regulation S-P.
- We will conduct post-incident review and share relevant findings where appropriate.
- We will honor data-rights requests (access, correction, deletion, portability) subject to our regulatory record-retention obligations.
- We will publish material changes to this posture on a forthcoming trust center (`trust.protocolwealthllc.c`) in addition to version-controlling this document.

Ongoing investment

- Quarterly internal security reviews
- Quarterly ZDR configuration audit (next: July 21, 2026)
- Quarterly tabletop incident response exercises (started Q2 2026)
- Annual vendor security re-reviews
- Periodic external penetration testing (formal engagement planned for 2026)
- SOC 2 Type I readiness targeted for Q4 2026; Type II attestation on a 2027 timeline
- BlockSkunk Phase 0/1 (ISO 27001 + SOC 2 alignment + GCP environment review) — engaged and underway

---

About this document

This document is the public-facing summary of Protocol Wealth's privacy, security, and compliance posture, captured on 2026-05-19. It complements:

- The partner-facing posture (`shared/docs/compliance/security-posture-partner.md`), which carries additional regulator-grade specifics under NDA-appropriate framing.
- The advisor-facing reference (`shared/docs/firm/advisor-reference-2026-05-19.md`), which covers operational rituals and in-product workflows for PW advisers.

For the most current state of in-production infrastructure, see `shared/strategy/CURRENT-STATE.md`. For the canonical compliance document set with PDF status, see `shared/docs/compliance/README.md`.

Protocol Wealth, LLC | SEC-Registered Investment Adviser | CRD #335298