p4-cache — Security & Compliance Brief

Foundational document for enterprise security reviews. Concrete defenses, citations to architecture decisions, third-party-style audit history.

Summary

p4-cache is purpose-built for environments where Perforce uptime is non-negotiable and audit trails matter. The product has been through three independent code audits (quality, security, and verify-binary verification). It ships with a documented threat model, ADR-tracked security decisions, gate-enforced supply-chain checks, fuzz-tested protocol parsers, and streaming-MD5-verified restore integrity against p4d's own `db.storage.digest`. The workspace is ~47K lines of Rust across six crates, with workspace-wide enforcement of `unsafe_op_in_unsafe_fn = "deny"` and `await_holding_lock = "deny"` — both compile-time defenses against entire classes of bugs. The runtime ships with a forensic watermark embedded into every cold-tier write so a leaked or exfiltrated object is traceable to the licensed deployment that produced it.

This brief is the document an enterprise security team can take to their CISO. It does not replace a formal security review of your specific deployment; it does answer the questions security reviews routinely ask.

Threat model summary

In scope:

- Malicious or compromised local user attempting to issue arbitrary trigger requests or pollute the audit trail
- Compromised cloud backend or MITM bypassing TLS validation attempting to inject corrupted content
- Operator misconfiguration (typo'd endpoint, missing TLS verification, world-readable credentials)
- Backend transient failures, queue overflow, or sink unavailability causing audit-log loss
- Post-build binary tampering (patched-out license check, NOP'd audit writes)
- Wall-clock manipulation across restarts to load a license that's past its grace window (containers, unprivileged user namespaces, virtualised time)
- `P4CACHE_LICENSE` env-var redirect pointing at an attacker-controlled file
- License-audit-log truncation to hide over-capacity transitions
- Signature-algorithm downgrade against the license envelope

Out of scope:

- Compromised root on the daemon host (root can do anything anyway)
- Compromise of the customer's Perforce credentials or `p4d` itself

- Physical access to the cold tier (durability and access control are the object store's responsibility — Azure Blob, S3, GCS each have their own well-documented models)

Defenses by category

Process-level isolation

Per-call trigger binary; no in-process libc hooks in p4d. The v2 integration is a per-call `p4cache-trigger` binary that p4d spawns via its `+X` archive trigger. The trigger speaks postcard frames over a local Unix socket (or named pipe on Windows) to the daemon and exits when done. A trigger crash exits non-zero and surfaces as a stock p4d archive-op failure; it cannot affect p4d's process state.

Daemon runs as an unprivileged user. Production systemd units run the daemon as `perforce`, not root.

Panic handling is `unwind`, not `abort`. The daemon's per-connection trigger handler is wrapped in `catch_unwind`; a panic during a single request becomes a structured ERROR log rather than killing the listener.

Authentication boundaries

Peer-credential gating on the trigger socket. On Unix, `SO_PEERCRED` is checked against the configured `trigger_allowed_uids` allowlist on every accepted connection. Peers outside the allowlist get the connection dropped before any protocol bytes are exchanged. On Windows the named-pipe peer-credentialing story is documented in `docs/windows-port-status.md`; today the listener fails-closed on any peer-check failure (including the Windows "Unsupported" verdict) rather than admit a connection silently.

Input validation at the trust boundary. Depot paths and revision strings are validated against traversal, NUL, LF/CR, oversize, and `-`-prefix argv tricks before any handler reaches the filesystem, the bucket key, or `p4` argv. The validators are shared between the trigger and the daemon so a compromised trigger or operator-misconfigured invocation cannot push junk through the socket.

Reserved depot subtree. No depot file may live under any `.p4cache` component at any path depth — enforced by the shared `validate_rel_path_bytes` validator and property-tested in `shared/src/lib.rs`.

Data integrity

Streaming MD5 against `db.storage.digest`, end-to-end. Every cold restore is hashed as bytes stream from the cloud backend into the cache temp file. Before the temp is atomically renamed onto the live archive path, the streaming digest is compared against `db.storage.digest` — the same MD5 p4d itself stores in `db.storage`. Mismatches fail the restore; rejected bytes never reach p4d.

Why MD5, not BLAKE3, on the restore path. p4d already maintains an MD5 for every archive. Re-using that digest as the verification anchor means a corrupted cold-tier object fails closed against p4d's own ground truth, without the daemon maintaining a parallel hash database that could drift. BLAKE3 still drives the tamper-evident license-audit keyed-MAC chain and the binary integrity self-check (the `dl_iterate_phdr` PT_LOAD hash baked in by the `integrity-stamp` tool) — neither of which is the restore install gate; that stays MD5 against `db.storage.digest`.

Atomic cache writes. Every write goes through temp + fsync + rename + fsync(parent). The path/container/rev layout matches p4d's exactly; the gzip encoding is content-based (incompressible payloads are stored raw — p4d's own unconditional gzip would grow them). Both regimes are verified by `rust/tests/p4_gzip_parity.sh`.

Endpoint pinning. Custom backend endpoints (private Azure endpoints, MinIO, GCS mocks) require explicit `endpoint_pin = true` in config. Without the opt-in, the loader refuses to start. Defense against credential-redirection attacks where a typo or env-var injection points uploads at a wrong host.

Forensic watermark on every cold-tier write. The runtime embeds a deployment-identifying watermark — derived from the Ed25519-signed license envelope's `customer_id` + `license_id` plus a per-process `daemon_instance_id` — into every cold-tier upload. The watermark surfaces as:

- Per-object custom metadata: `x-amz-meta-*` on S3, `x-ms-meta-*` on Azure, `x-goog-meta-*` on GCS (carrying `customer_id` / `license_id` / `daemon_instance_id` / `source` / `tier`). Visible to anyone with backend-read access; that's the design intent (a leaked or exfiltrated blob is traceable). This is the live forensic channel.
- The SDK `User-Agent` is set to the deployment suffix on every cloud request (via `object_store` `ClientOptions`), so it appears in AWS CloudTrail / Azure Activity Log; it is also surfaced in logs / Prometheus `license_info` at startup. The per-object metadata above is the second, log-independent channel.

Operationally passive — the watermark does not affect throughput, semantics, or correctness. Disabling it requires a custom build or an empty-watermark envelope; neither is shipped to customers. The anonymous `Watermark::empty()` exists only for the integration-test corpus so test artifacts don't carry watermarks.

License envelope and ceilings. The license envelope (Ed25519-signed RFC 8785 JCS-canonicalized JSON) carries a capacity claim and an expiry. The runtime drives a four-state capacity machine off the ratio of managed bytes to the soft ceiling (`capacity_tib × 1 TiB`):

State	Range	Behavior
Green	< 80%	normal
Warn	80–100%	warn log, serve everything
Grace	100–120%	SEV-1 log, serve everything
Refused	≥ 120%	refuse new admissions only; cold-tier uploads keep running (draining the cache is the only recovery path, so gating uploads on capacity would wedge the cache); reads and restores of already-tracked files continue

All four states surface via Prometheus (`p4cache_license_managed_bytes`, `p4cache_license_soft_ceiling_bytes`, `p4cache_license_hard_ceiling_bytes`) so operator alerting can fire long before Refused. State transitions are atomic — capacity and expiry pack into a single `AtomicU16` swap so the hot-path admission guard never blocks.

Algorithm pinning. `VendorPubKey` carries an explicit `signature_alg` field; the verifier requires both the envelope's declared algorithm and the matched key's algorithm to be `Ed25519` (verified with `verify_strict`). Key-id selection is an ordinary equality match against the embedded vendor-key table — the key id is a public, non-secret identifier. Signature parsing is explicit-length-checked rather than infallible — wrong-size signatures get `LicenseError::SignatureLength`; the code uses an explicit length check plus `copy_from_slice` (no `try_into`).

Past-expiry claims rejected at load time. Startup refuses a license whose `expires_at` is more than `EXPIRY_GRACE_DAYS` (30 days) past the current wall clock, backed by a cross-restart clock anchor, so a stale envelope can't quietly load. Capacity values ≥ 32 PiB (`1<<15` TiB) are also rejected at parse time (overflow defense).

Binary self-verification — hard-fail at startup. The `integrity-stamp` vendor tool digests every shipped artifact and bakes the value into a dedicated block in `.rodata`. At startup the daemon walks its own program headers via `dl_iterate_phdr`, hashes every `PT_LOAD` segment with the `PF_X` (executable) flag set in `p_vaddr` order, and refuses to start on mismatch. Hashing the loaded segments (not `/proc/self/exe`) catches in-memory patches; hashing every executable LOAD (not just `.text`) catches LTO/PGO/`-ffunction-sections` split-text patches that would slip past a single-section check. An opt-out (`P4CACHE_INTEGRITY_SOFT_FAIL=1`) exists for build-pipeline use and is loudly logged + audit-logged when honored.

License-load env hardening. The `P4CACHE_LICENSE` env var still lets operators point the daemon at a non-default license path, but the loader now refuses the override unless the target file is owned by the daemon's effective uid (or root), carries no group-write or world access (0600 or 0640), and is not a symlink. Every honored override emits an audit-grade tracing event at startup with the resolved path.

Cross-restart clock-skew anchor. The first successful license load persists the wall-clock view of "now" to `<state_dir>/license/clock-anchor.json`. Every subsequent startup compares the current clock to the anchor and refuses to load if the clock has moved backward by more than `CLOCK_SKEW_TOLERANCE_SECS` (1 hour — absorbs NTP corrections, leap-second-style backward steps, and DST mis-handling without false-positive). Closes the "rewind the clock to load an expired license" path that virtualised-time environments otherwise admit.

Tamper-evident license audit log. Every license-relevant event (load, capacity-state transition, expiry transition, integrity-check result) appends one NDJSON line to `<state_dir>/license-audit.log` with a BLAKE3 keyed-MAC chain. The MAC key is the per-deployment `customer_secret`. The format includes monotonic `seq` numbers, an RFC 8785 canonical HMAC input (via `serde_jcs`, so nested-object payloads can't produce ambiguous bytes), and a cross-restart anchor file (`license-audit.anchor.json`) recording `(latest_seq, latest_hmac)`. On startup the loader compares the anchor to the file's max seq and warns on mismatch — closing the "truncate the tail to hide bad events while keeping the chain internally valid" path. Developer-build licenses (all-zero `customer_secret`) switch to plaintext "unchained" lines with a `# DEVELOPER BUILD — UNCHAINED` header so operators don't mistake them for production audit history.

Transport security

Unified TLS stack: rustls everywhere. Postgres TLS was migrated from native-tls/OpenSSL to rustls (tokio-postgres-rustls) so the binary has one TLS stack with one CA-parsing surface. Same stack used by Azure, S3, GCS, and Postgres.

Private CA support for the audit-log sinks; system-trust-store for the cloud backends. The PostgreSQL and Elasticsearch audit sinks honor an explicit CA via config (`access_pg_ca_cert` / `access_es_ca_cert`), validated by `validate_ca_cert_file` before runtime use. The v2 cloud object-store backends (S3, Azure, GCS) use the host system trust store via their rustls client; their `ca_cert_path` config flag is **rejected fail-closed at config load** today — the daemon refuses rather than silently ignore an unsupported TLS option. For a private-CA cloud endpoint, install the CA in the host trust store. Explicit per-backend TLS-option wiring is tracked in the v2 architecture plan.

`verify_ssl = false` is rejected by the runtime for the cloud object-store backends — the daemon fails closed at config load rather than silently fall back to the system trust store against operator intent. A self-signed or private-CA endpoint is supported by installing its CA in the host system trust store (the cloud-backend `ca_cert_path` config flag is likewise rejected today, pending explicit TLS-option support).

Audit trail

Access logging to Elasticsearch or PostgreSQL. Every depot file read is recorded with path, size, and last-access timestamp. The pipeline:

- The trigger listener emits access events directly into an in-process channel as it serves each archive — the sole producer in the pipeline (no datagram socket is bound)
- Daemon deduplicates per batch (configurable batch size, default 10K events)
- Writer task flushes to Elasticsearch (TLS with custom CA, basic-auth) or Postgres (TLS with custom CA, SQLSTATE-aware error classification, automatic table creation with a validated identifier)
- Exponential backoff on remote failures (initial/max configurable)
- Auto-disable after configurable consecutive failures, with warning rate limit
- Optional on-disk spool (bounded by `access_spool_max_bytes`) for outage resilience; spool replayed oldest-first when the sink recovers

Drop boundaries are explicit and bounded. Two distinct loss points:

- Pre-spool: the bounded in-process channel between the trigger listener and the writer task. When the writer is behind, the producer `try_send`s and sheds the event rather than awaiting, so the audit pipeline can never apply backpressure to the hot serve path. Operators size this channel with `access_max_pending_entries`.
- Post-spool: when the on-disk spool reaches `access_spool_max_bytes`, new undeliverable batches are dropped, and invalid/unreadable spool entries are surfaced via the `p4cache_access_spool_invalid_dropped_total` Prometheus counter.

The serve path never blocks on audit delivery: pre-spool overflow is a deliberate, bounded shed governed by `access_max_pending_entries`; the spool absorbs sink outages, and its drop surface is named, configurable, and observable in Prometheus.

Supply chain

Gate-enforced checks (per `scripts/gate.sh`, run on every change):

- `cargo audit` on every gated change
- `cargo deny check advisories bans licenses sources` on every gated change
- License allowlist enforced (no GPL-only deps; permissive-license policy documented in `rust/deny.toml`)

- The same gate runs natively on the Windows dev machine: daemon/trigger build, the cfg(windows) clippy wall, and the Windows unit + integration test set

Cargo dependency hygiene. Direct deps are reviewed; `cargo audit` and `cargo deny` both run green with zero vulnerabilities, and cargo-deny's active ignore list is now empty. The only accepted exceptions are two *unmaintained*-status warnings (not vulnerabilities), both transitive/dev, documented with reachability rationales in `rust/.cargo/audit.toml`: `rustls-pemfile` (RUSTSEC-2025-0134, now pulled only transitively by `object_store` — the daemon's own direct use was replaced with `rustls-pki-types`) and `bincode` (RUSTSEC-2025-0141, dev-only via `iai-callgrind`). Workspace-wide lints enforce `unsafe_op_in_unsafe_fn = "deny"` (Rust 2024 idiom — every unsafe operation inside an `unsafe fn` body requires an explicit `unsafe { ... }` block) and `await_holding_lock = "deny"` (compile-time prevention of guard-across-await deadlocks). Silent integer casts are denied under `cast_sign_loss = "deny"`, `cast_possible_truncation = "deny"`, and `cast_possible_wrap = "deny"`; surviving sites use the `wrapping_as_*` helpers (intentional wrap) or `try_from(x).expect("…")` (panic-on-bug at a boundary).

Rust 2024 edition. The workspace is pinned to edition 2024 with MSRV 1.88; the gate builds on exactly that toolchain so a change accidentally using a newer language feature gets caught.

Fuzz testing

Two `cargo fuzz` targets run in the full gate's smoke mode:

- `fuzz_handle_request` — the daemon's trigger request handler
- `fuzz_checkpoint_parser` — the Perforce checkpoint journal parser

Both run ~60 seconds per CI invocation; longer runs can be invoked locally with `cargo +nightly fuzz run`.

Memory safety

Every unsafe block in the daemon carries a `// SAFETY:` comment justifying the invariants. The unsafe surface covers `dl_iterate_phdr` self-hash over executable LOAD segments, cross-restart clock anchor, license audit-log v2, SCM_RIGHTS fd-passing on the trigger socket, and the checkpoint-verifier's libc/POSIX integrations; every block is reviewed under the same SAFETY-comment policy.

Workspace-wide lints `cast_sign_loss = "deny"`, `cast_possible_truncation = "deny"`, and `cast_possible_wrap = "deny"` catch silent integer-cast bugs at compile time; surviving sites either go through the `wrapping_as_*` helpers in `p4cache_shared` (intentional wrap) or `u32::try_from(x).expect(...)` (panic-on-bug-at-boundary).

Audit history

p4-cache has been through three independent code audits, each with synthesis reports tracked in the repo:

Audit	Scope	Outcome
`.audit/`	Quality, architecture, security, performance, tech debt — full Rust workspace	80+ findings across phases; all CRITICAL findings closed; HIGH and MEDIUM remediation tracked in

		`.audit/6-synthesis/REPORT.md`
`.codex-audit/`	Revalidation against latest audit-remediation commits, deeper storage-SDK retry/perf review, docs and packaging drift	All HIGH engineering findings (e.g. `S-004` staging-path RAII) closed with pinning tests. Two supply-chain HIGHs (`S-001`, `S-003`) flagged against upstream SDKs have since been cleared (lockfile `cargo update` plus dropping the unused `protobuf` feature); `cargo audit`/`cargo deny` now run green, with only two *unmaintained*-status warnings accepted — `rustls-pemfile` (transitive via `object_store`) and `bincode` (dev-only).
`.audit-verify/`	Targeted verification audit of the checkpoint verifier (now the `p4-cache verify` subcommand)	One HIGH (argv-injection hardening), small-scope MEDIUMs around resource limits and logging unification

Health scorecard from `.audit/6-synthesis/REPORT.md` (re-graded 2026-05-31 against the current v2.1.0 tree across successive in-repo lift passes, with adversarial anti-inflation calibration — each grade independently verified against the live tree, e.g. the perf gate was confirmed to fail on an injected regression):

Axis	Rating	Notes
Code quality	A	`clippy::pedantic` + cast-safety (`cast_possible_truncation` / `cast_sign_loss` / `cast_possible_wrap` = deny) + `undocumented_unsafe_blocks` + **`unwrap_used` + `expect_used` = deny**, all enforced under CI `-D warnings` — no unguarded panic can land in a production path without a justified `#[allow(reason = …)]`; zero `TODO`/`FIXME`, zero `#[ignore]` tests; typed `thiserror` errors throughout; **80.5% core line coverage behind a CI floor gate** (every translation-layer module ≥ 85%; full-workspace line coverage 58.1% — the "core" figure excludes integration-only surfaces, see `docs/perf/coverage.md`). Held from A+ by core-not-full

		coverage and the absence of an external SAST/code-quality attestation.
Architecture	A−	Acyclic 6-crate workspace; daemon/trigger split; one production storage trait impl behind a factored atomic-download primitive; the v1 god-object/shim fully retired into focused subsystems; dead forward-compat trait surface removed; ADRs for the major decisions. Held from A by a few large (cohesive) modules not yet sub-divided.
Security	A	All 32 security findings closed; every `unsafe` block carries a `// SAFETY:` comment (deny-enforced), and `unwrap_used`/`expect_used` are `deny` so no unguarded panic ships; restore-time MD5 against `db.storage.digest` (fail-closed); forensic watermark on every cold-tier write; `dl_iterate_phdr` self-integrity check (refuses to run on mismatch); hardened license path (cross-restart clock anchor, owner-or-root + 0600/0640 env-file gating, algorithm pinning, tamper-evident BLAKE3-MAC audit log); fail-closed peer-cred and config validation; rustls everywhere; 2 `cargo fuzz` targets in the gate. **`cargo audit` + `cargo deny` are green with zero actual vulnerabilities** — the HIGH `quinn-proto` QUIC-DoS and the `protobuf` uncontrolled-recursion advisories were both eliminated; the only accepted exceptions are two *unmaintained*-status warnings (not vulnerabilities) — `rustls-pemfile` (transitive via `object_store`) and `bincode` (dev-only via `iai-callgrind`) — and cargo-deny's active ignore list is now empty, with the supply chain now self-policing (Dependabot, a stale-ignore

		hygiene check in the full gate, signed-release build provenance, `SECURITY.md`). Held from A+ by the absence of *independent external* validation — a third-party penetration test / SOC 2 (see [What this brief is not](#what-this-brief-is-not)) — an A→A+ differentiator, not an internal-posture gap.
Performance	A−	Five criterion benches with baselines in `docs/perf/baseline.md`; an end-to-end read-stress harness; `p4 -ztag storage` taken off the read hot path (v2.1.0); single-pass streaming restore hash; thin-LTO + strip (and dropping an unused `protobuf` exposition feature) cut the daemon from 52 MiB to 13.4 MiB, budget-gated. **A deterministic `iai-callgrind` instruction-count gate now fails the build on a >10% per-op regression.** Held from A because the deterministic gate covers per-op CPU floors, not the end-to-end restore/read paths (still hand-captured wall-clock numbers).
Tech debt	A	Zero `TODO`/`FIXME`/`HACK`; ADRs 0001–0003 + CHANGELOG; deferred control-flow lints converted and now enforced; no-unguarded-panic machine-enforced; coverage-floor backlog empty; MSRV + supply-chain gates; the one known flaky test was root-cause fixed. In-repo debt the team can itself discharge is essentially zero — the only residual is two *unmaintained*-status advisories (un-fixable here): `rustls-pemfile` (transitive via `object_store`) and `bincode` (dev-only via `iai-callgrind`), plus the absence of an external attestation, which caps it from A+.

No CRITICAL findings open today. The audit framework defines CRITICAL as "data loss, RCE, auth bypass, secret leak, persisted state corruption, or known-exploit CVE in a direct dep" — no item in the workspace meets that bar.

Audit infrastructure itself — three independent passes with tracked remediation — is unusual for software at this maturity. Most year-1 commercial products have zero audit trail.

Compliance posture

p4-cache does not currently hold formal compliance certifications (SOC 2, ISO 27001, FedRAMP). The product *enables* customers to meet their own compliance requirements through:

- **Auditable access logs** for every depot file read, sinkable to the customer's own SIEM (Elasticsearch) or audit warehouse (PostgreSQL). Standard SOC 2 / HIPAA / FedRAMP access-logging requirement.
- **Encryption in transit** to all backends via rustls.
- **Encryption at rest** delegated to the backend (Azure Storage Service Encryption, AWS S3 server-side encryption, GCS default encryption, NFS at customer discretion). Server-side encryption is configured on the bucket's own default-encryption policy on the provider side; the daemon **rejects** a `server_side_encryption = true` config flag at load (fail-closed) rather than imply an SSE guarantee it does not itself enforce.
- **No customer data crosses the network in cleartext.** TLS everywhere (rustls). For air-gapped or private-CA environments, the audit sinks accept a CA via config and the cloud backends use the host system trust store.
- **No telemetry, no phone-home.** The daemon does not contact any external service except the customer-configured backends. Suitable for air-gapped and high-classification environments.

For customers requiring formal compliance, p4-cache can be deployed on infrastructure that holds the certifications (AWS GovCloud, Azure Government, customer-controlled on-prem). The product itself is a software component on that infrastructure; the customer's existing compliance boundary covers it.

Reference customer / case study posture

Currently pre-production; reference customer availability is limited until the first two enterprise deployments complete their first quarter in production. Design-partner-tier customers (early adopters with reference rights agreements) are available now under NDA.

What this brief is not

It is not a substitute for:

- A formal third-party penetration test on your specific deployment (recommended for any production rollout in regulated industries)
- A formal SOC 2 audit report (which p4-cache does not currently hold)

- A signed Master Services Agreement and Data Processing Addendum, which become part of the commercial relationship

It is a *starting point* for your security team's review — enough to know what you'd be deploying, what defenses are in place, and what questions to ask. Most security teams ask for the architecture deep-dive next; that's available as a separate document.

Document control

- Version: 1.0
- Maintained against repo HEAD; refresh on each substantive security change
- Citations to specific ADRs and audit reports indicate verifiable claims, not marketing language
- Subject to NDA for the audit-finding-level detail; this public brief lives at the summary level by design