STANDARD

PUBLIC SAFETY EDITION

BY METRASLATE

THE DEFENSIBLE RECORD AROUND THE EMERGENCY CALL PATH

PRODUCT DOCUMENT
SEPTEMBER 2026

METRASLATE LLC

1. Summary

An Emergency Communications Center buys its call path in pieces — the ESInet that delivers the call, the console that answers it, the CAD that dispatches it — and runs the organization around that path on a scheduling product, a QA product, a training-records product, a records-request tracker, an equipment log, spreadsheets, and paper. Each of those products is competent at its job. None of them holds the incident, none of them holds each other's history, and when litigation, a records request, or an after-action review asks what the telecommunicator knew and when, the answer is reconstructed by joining vendor logs on clocks that were never reconciled.

Standard is a domain-agnostic, multi-tenant work platform that ships as editions. It proved itself first in the commercial calibration industry, where a record produced today must be reproducible and defensible for the life of the instrument, and the substrate that industry forced — its record spine — append-only records, hash-verified document reproduction, database-enforced fail-closed tenancy — carries into every edition unchanged, along with its operations layer. What the edition adds is a record of the incident built to four commitments:

1. **It preserves what was knowable** — by whom, from what source, at what time — and never lets later knowledge rewrite an earlier decision.
2. **It refuses what it has no business knowing.** Every element it holds can state why it is there; what belongs to CAD, the recorder, the RMS, the ePCR, or the state system stays there, referenced and not possessed.
3. **It never turns a guess into a fact.** Every statement it makes is graded — verified, received, derived, asserted, or absent — and a model's output is an inference with an interval, never a label.
4. **It calibrates its own instruments.** Every model whose output enters the record has a calibration record, reference cases, a validity domain, and a kill switch, so that when a person appears to have changed, the record can show that the measurement did not.

The edition does not replace the ESInet, the console, or CAD. It integrates with them, captures what they emit as a single tamper-evident history per agency, and runs the agency on top of that history: who is on shift, which positions and units are covered, which incidents need review, which equipment is due, which records request is at deadline, and how the center performs against its region. Its value is independence: it is the one record in the building that does not belong to a call-path vendor, and it is built to audit the vendors' own systems — including their AI — rather than to be one.

It is built by Metra Praxis: the person who knows the work describes it in her own words; the description becomes a written picture of the whole system; the picture is grilled and ratified; the first product is one pain removed in thirty days; and nothing is built from a guess. The record names in this document — occurrence, thread, link, variance, finding, knowledge state — are working names, hers to replace.

2. The problem

2.1 The call path is solved; the organization is not

A modern center receives a call over an NG911 ESInet, answers it on a call-handling platform, receives device-derived location, medical profile, crash telemetry, and video through a supplemental data broker, and creates and manages the incident in CAD. Each layer has entrenched vendors and multi-year procurement

cycles. Around that call path, the agency schedules people to console positions, tracks their certifications, reviews their calls, answers public-records requests on statutory deadlines, produces court packets, reports to the council, and coordinates with neighboring agencies. The fire and EMS agencies it dispatches maintain apparatus and life-safety equipment on mandatory inspection and calibration intervals. Every one of those tasks has a product. The products do not share a record.

2.2 The record that matters is the reconstruction

Incident history is scattered across CAD exports, the logging recorder, broker portals, the protocol system, and paper. Logging recorders reconstruct audio, screen, CAD, and radio on one timeline, but the reconstruction lives in the recorder vendor's product and stops at that vendor's inputs. As AI assistance enters the console — transcription, translation, summarization, protocol suggestion — no independent record holds what the machine suggested beside what the human did, and no one is checking whether the machine's judgment drifts.

The record whose being wrong is catastrophic is not the QA form or the certification file. It is the reconstructed timeline: what was received, what each system said, what each person could have known at each moment, what protocol recommended, what policy required, what the human decided, what the field did, and where those diverged. That record is litigated, subpoenaed, and reviewed after every critical incident, and today it does not exist as a record. It exists as a reconstruction someone performs.

2.3 The costs

- **Litigation and discovery.** A defensible timeline today means subpoenaing three vendors and reconciling three clocks by hand.
- **Quality assurance.** Reviewers look backward with information the telecommunicator could not have had, and sampling and routing are done by hand.
- **Public-records requests.** Manual, slow, inconsistent, and unable to prove what was redacted and why.
- **Staffing.** Schedules live outside the system that knows the call volume.
- **Consolidation.** Regional centers absorb agencies whose history arrives as spreadsheets.
- **Equipment.** Life-safety devices are tracked with no connection to the incidents they were used on.
- **AI.** Vendor models enter the console with no independent record of their versions, their limits, or their drift.

3. The platform

Standard is one platform that ships as editions: the substrate — records, tenancy, documents, operations, numerics — is shared, and an edition declares its subjects, vocabulary, numbering, document templates, and reports in a manifest. The substrate carries no assumption about the industry it records; calibration was the first edition and the proving ground. The following are built.

The record spine. Records are never updated in place; a correction is a new record referencing the one it supersedes; the full history of any object is a read. The document engine is welded to the records that feed it: a rendered document carries a hash of its source data and a deterministic vocabulary projection, so a re-render years later either reproduces it exactly or fails verification. The value path carries exact decimals with explicit scale.

The tenancy foundation. Every tenant-scoped table carries row-level security enabled and forced; a connection without tenant context returns zero rows. Large tenants run on dedicated databases, the long tail shares a schema under RLS, and both resolve through a placement layer so a tenant can move tiers without code changes. Tenant context propagates through asynchronous execution. All indexes are tenant-leading. Standard has already found and closed an existence-and-count oracle in free-text search under RLS; the discipline that closed it governs every classified column in this edition.

The operations layer. Dispatch, role dashboards, the pipeline board, schedule and timecard, My Metrics, the persistent dock, personnel records, customer identity and book of business, in-context chat, the workflow engine, and the asset spine with work orders.

Specified, next to build. The domain-event emission layer, which the edition's ingest adapters attach to, and systems-integrity monitoring, which turns integration liveness and platform health into an operational signal. Both are decided in the platform's decision record and compose primitives that exist; both are on the edition's critical path.

The method. Standard is built by Praxis: briefs, adversarial grills, lanes, refuters, fences with planted violations watched to fail, receipts, and landings on a recorded word. Every claim in the process carries a tag — observed, claimed, inferred. The edition's record inherits that epistemology directly (Section 4.9).

4. The record model

4.1 Occurrence, thread, event, link

The real-world event is an **occurrence**. Each source system's own record of it is a **thread**, kept whole and unmerged: a call leg, a protocol session, a CAD incident, each additional caller, a unit's response, and later an RMS case, a fire report, an ePCR. Every atomic fact from a source is an **event** on a thread. A **link** is an assertion that two threads, or two occurrences, are related. Standard never decides that two things are the same event; it records who said so, on what basis, at what confidence class.

The link ladder. *Candidate* — the system, from a deterministic rule the agency configured, with the evidence stated; merges nothing and is not shown at a live position unless the agency chose that and the presentation is captured. *Asserted* — a named person or an authoritative source's own function (CAD's associate or duplicate). *Confirmed* — a second, independent basis. *Disproven* — a named person with the reason; the link is superseded, every view that used it is re-derived, nothing earlier is removed. "We don't know yet" is an unlinked thread, a legitimate state.

The relation set. `same-occurrence`, `associated`, `derivative`, `split`, `merge`, `supersedes`, `duplicate`, each directional or symmetric, each climbing the ladder independently. The occurrence is not a container; two occurrences can be associated without becoming one. A traffic stop that may involve the suspect from a domestic is its own occurrence, associated with the domestic through case pointers; the person is never held (Section 4.8).

Classification is a series. Each thread carries classification assertions as they stood at each change — CAD type code, protocol determinant, disposition — and "current classification" is a projection for a chosen time. Protocol transitions are paired events on both sessions with the reason and the initiator.

Custody state. Every thread and event is *held* (Standard holds a copy with provenance and hash), *received* (another custodian sent it; attributed to them; never nominated as authoritative), or *referenced* (a pointer: the

record exists, its custodian is named, its content is not here). A referenced thread is a real row, so "their authoritative record is unavailable to us" is a statement, not a hole. A blast-radius traversal stops at a referenced node and reports the boundary and the count it could not reach. Where the neighboring agency is another tenant, a reference can become a cross-tenant link by both tenants' recorded consent, under both tenants' row-level security.

4.2 Four timestamps and the knowledge state

Every event carries a *source time*, an *ingest time*, a *reconciled time* with an uncertainty interval (Section 4.5), and, for anything shown to a person, a *presentation time* per principal. Presentation is its own event family — rendered, acknowledged — with a capture grade: **acknowledged** (the source reported the person interacted with the item), **rendered** (the source reported it was on screen), or **available** (it had reached the position; presentation not captured; inferred from arrival and labeled so).

The **knowledge state** of a principal at a time is the set of events presented to that principal by then. A decision event references the knowledge state it was made in. Review opens in **as-known mode** by default: only what had been presented to the chosen principal before the chosen decision, with later events in a separate lane; switching to hindsight is recorded on the review. Relays are events with chains, so what an officer told the dispatcher and what the dispatcher told the caller are two events joined by a basis reference.

Presentation evidence comes from three sources in descending grade: the recorder's screen capture, which lets a reviewer raise a presentation to *verified* by confirming an item on the captured screen at the moment of decision; Standard's own panels, where Standard is the presenter of a feed and emits render and acknowledge events natively (whether a Standard panel at the position crosses the off-limits line is the Director's ruling); and vendor telemetry where the console or CAD emits it.

4.3 Basis, authority, and the decision chain

Every decision event carries a **basis** (the events it consumed, at a stated capture grade: explicit, by presentation, or by availability) and an **authority** — protocol, agency policy with the version cited, system recommendation, telecommunicator, field command with rank and role, external agency — with the actor and the policy citation where one applies.

The dispatch decision is five events, never merged: protocol determinant; policy response, resolved by Standard against the agency's **response plan**, which lives in Standard as a versioned configuration record so that "what policy required" is answerable against the plan as it stood; CAD recommendation, from the CAD feed or parsed from the CAD chronology where the CAD writes it there; dispatch executed, by the dispatcher; field modification, by field command or relayed. Divergence between layers is a variance written at ingest (Section 7.2). A layer the source does not expose is recorded *absent*, so the chain shows a gap rather than a silent skip. Lineage runs through basis references in both directions: from a translation forward to every decision that could have rested on it, and from a dispatch back to the audio segment.

4.4 Location assertions

There is no location field. A **location assertion** is a record family with one row per assertion: source (carrier ALI or NG911 location object, device-based supplemental location, caller-stated, CAD geoverification, prior call history, responder-confirmed); the source's own time and Standard's; geometry as the source gave it — address, point with radius, polygon, floor — never collapsed; the source's own confidence or accuracy where provided and *not provided* where not; the purpose the source serves in the agency's practice; and

presentation per position. A better location is a new assertion; a confirmed location is a new assertion whose basis references the one it confirms or refutes.

"What location did the telecommunicator reasonably believe was actionable at 14:03:17, and why" is a projection over the assertions in that person's knowledge state at that instant, the agency's **location policy** as it stood (a versioned configuration record of source precedence), and the basis of the decision made then. The responder-confirmed location arrives later with its own time, does not exist in as-known mode at 14:03:17, and a finding that cites it against that decision is rejected at the writer.

4.5 The authority register and the clock model

Standard is authoritative for exactly two things: the fact of receipt (what arrived, from where, when, in what form, with what hash) and the assertions made inside it. For everything else it is a custodian of copies and says so.

The **authority register** is the agency's versioned declaration of which source owns each data element, across the boundary as well: for every element about another agency's incident, the authoritative source is that agency and Standard's copy is never nominated. Every source's value for an element is kept; a projection uses the nominated one; disagreement outside the sources' uncertainty is a **discrepancy** record with a lifecycle, explained by an assertion with an author, never resolved silently.

Every source has a **clock model** — offset and drift against Standard's reference, with an uncertainty — fitted from two streams: event pairs that should coincide, and the source's own reported time sampled on every exchange by the adapter. Reconciled time is an interval; overlapping intervals are reported as indeterminate order rather than forced. This is a measurement statement — a value, a correction, an uncertainty — and it is the same discipline Standard applies to a calibration reading.

4.6 The classification register and least privilege

A second agency register says what each element *is* — criminal-justice information, protected health information, personally identifying, public — and which principals, roles, external processors, and administrators may read each class. Authorization is enforced at the database in the same fail-closed row-level security that isolates tenants, extended with class predicates. There is no role that sees everything; an administrator's role is class-scoped; break-glass is a recorded, alarmed, time-boxed grant the integrity monitor reports.

Existence is classified. A link inherits the higher classification of its endpoints. Everything a principal can see — nodes, edges, counts, aggregates, search results, blast-radius outputs, reports — is computed inside that principal's authorization. There is no "3 hidden items" and no edge to a node one cannot open. Only leakproof predicates on classified columns; no error-message or timing differentials; no counts computed outside the policy.

Audit is not a second repository. Audit events carry identifiers and actions — who, which record identifier, which action, when, from where, under what authorization — never payloads, never field values, never a search string. An audit event inherits the classification of what it references. The audit log is append-only, retained by class, and its own accesses are audited.

External processors are principals. AI services, brokers, and vendors sit in the classification register like employees. A class may be sent to a processor only where the register permits, and every dispatch is an audit event with the class named.

4.7 Commitments, retention, and lawful destruction

Immutability means no silent mutation. It does not mean eternal retention.

Commitments. Every event is stored as an *envelope* (identifiers, times, source, class, size, custody state, the commitment) and a *payload* (content-addressed, encrypted under a data key). The commitment is a hash over the payload and a random nonce that lives only with the payload. While the payload exists, verification includes the nonce; after destruction, the retained commitment proves that a record existed with that identifier at that time and cannot be brute-forced against guesses of low-entropy content. The hash chain that seals Standard's records verifies over the surviving structure. The same commitment lets an authorized party later prove that a classified link existed at time T, while an unauthorized principal never sees the link or a gap where it was.

Four operations. *Retention*: incident events are partitioned by tenant, month, and record class, because statutory periods differ by class within one incident; retention is a partition drop preceded by hash-verified extraction of any held or unexpired incident's events, plus destruction of the partition's data key. *Legal hold*: the key is retained, the partition preserved past schedule, and held media is copied in from the recorder (Section 5.4). *Restriction and sealing*: authorization events that narrow the readers, content retained, the order referenced. *Expungement and destruction*: a destruction event; payload and nonce destroyed; the envelope retains a certificate with the order, the rule version, the date, and the actor.

Derived records and links. Because every derived artifact carries basis references, destruction is a reverse traversal: everything derived from a destroyed event is destroyed or re-derived without it, with the certificate cited. A link's envelope survives unless the order covers the relationship, in which case both endpoints carry the certificate. A system without lineage cannot honestly claim to have destroyed anything.

Backups and proof. Data keys are per tenant, per class, per period under a managed key hierarchy; backups hold ciphertext; destroying a key renders every backup copy unreadable at once, and the window during which unreadable ciphertext persists is stated in days. Deliveries already made are listed on the certificate so recipients can be notified. Destruction is proven by a walked restore drill: a planted destruction, then an attempted read through every path including a backup restored to a scratch cluster, each of which must fail, with the failures as the receipt.

4.8 The four decisions: ingest, reference, derive, never touch

Standard possesses what it must be able to reproduce, references what someone else must produce, derives what it can recompute, and never touches what it has no lawful purpose to hold. Applied per element in the authority register, in this order; the first test that fails decides.

Never touch when the element is the call path itself (audio in transit, routing, call control, dispatch commands); the content of a criminal-justice query or its response (Standard may hold the fact that CAD logged a query, with its identifier and class, never what came back); a person's identity resolved across incidents (an RMS function); a government identifier; clinical content beyond the pointer; anything a data provider's terms forbid; or anything whose only justification is that it was available. The register has no content row for it and the adapter cannot write its payload.

Reference when an authoritative custodian exists with retention at least as long as Standard's need and Standard's purpose is to locate or verify, not reproduce: audio, video, and screen recordings; RMS case content; ePCR content; other agencies' records; raw supplemental payloads under restrictive terms. A pointer, a custodian, a hash where exposed, and the date established.

Ingest the payload when Standard is the only place the joined fact can be proven (receipt, timing, presentation), when the element is needed to reproduce a document Standard seals, when the custodian's retention is shorter than the need, and when there is a lawful basis and the terms permit. The envelope of every event is always ingested; the payload only under a stated condition on the register row.

Derive everything that is a function of what is held — knowledge state, reconciled chronology, variances, context packets, blast radius, reconstruction reports, reference renderings — stored with lineage, graded derived, disposable, and never nominated as authoritative.

A **data-minimization check** runs on a schedule inside systems-integrity monitoring: any held payload whose register row carries no ingest condition, or whose condition has lapsed, is reported for destruction or re-justification. Standard therefore does not know who the caller is, who the suspect is, what the state system returned, what was said on the tactical channel, or what the medic did. It knows that each has a record, where it is, who holds it, and what its own telecommunicators were shown and did.

4.9 Evidence grades

Every statement the platform makes carries one grade: **verified** (a copy whose hash matches the source's integrity value or which the source re-confirmed), **received** (a copy as received, the source providing no integrity value), **derived** (Standard's own computation, with lineage), **asserted** (a human statement made in Standard, with its author), or **absent** (expected under the register or the adapter's contract and not present — including *absent with pointer* for referenced custody and *absent with reason* for a refused inference). Absent is a value. The grades are the method's own tags — observed, claimed, inferred — applied to incidents instead of briefs.

5. Ingest and integration

Every adapter writes to a durable, replayable queue before anything touches the database; every write is idempotent on the source's own identifier; every event carries provenance — source system, source record identifier, adapter version, payload hash. Payloads the adapter cannot classify are stored raw under an explicit *unrecognized* class and surfaced by integrity monitoring; a logger that drops what it does not understand is not a logger.

5.1 Supplemental data brokers

RapidSOS is the primary integration; Rave, alarm-industry feeds, and other brokers follow the same pattern into a broker-agnostic schema. Raw payloads are held where the provider's terms permit and referenced by hash otherwise; the terms are verified before the adapter ships; medical profile data is handled under the PHI class.

5.2 CAD and call handling

Incident, unit-status, and chronology feeds from the major CAD platforms; chronology parsing for recommendation lines where the CAD writes them; EIDO for CAD-to-CAD, held with custody *received* and attributed. Console event interfaces are proprietary; the standardized path is NENA i3 logging.

5.3 NENA i3 logging — LogEvent consumer

The edition implements the i3 Logging Service interface for LogEvents from ESInet and call-handling functional elements. It does not capture SIPREC media; the recorder remains the media store, linked by call

identifier. Until conformance is complete , the edition is a records platform that also receives log events, not a certified NG911 logger.

5.4 The recorder

Audio, screen, and radio are referenced by segment. A legal hold copies the held occurrence's segments into Standard's content-addressed store, hash on entry, so held media is *verified* and under Standard's retention; integrity monitoring alarms on a hold whose copy has not completed within the agency's window, and recorders without an export interface get a custodian's task with the same alarm. Screen capture is indexed by decision for as-known review. Feature extraction for any model runs against the recorder's segment under the applicable scope, and only the features come back.

5.5 AI lineage

Every AI artifact — language identification, transcript, translation, summary, protocol suggestion — is a chain: audio segment reference; the artifact with model identity, version, configuration hash, confidence where exposed and *not provided* where not; the presentation of the artifact to the position with capture grade; the human decision with basis references; and any later human correction as a separate event referencing the artifact it corrects, marked post-hoc. The correction never touches the artifact. A reviewer can ask which decisions could have rested on a corrected translation, and the answer states the basis grade it was computed at. Every artifact references the calibration record of the model that produced it (Section 6).

5.6 Reference renderings

Where a vendor exposes no model metadata, Standard produces a reference rendering at review time — a second transcript and translation of the held segment by a model Standard controls, with its own version and calibration record — graded derived, never replacing what the call-taker saw, and disagreement between vendor and reference is a discrepancy record. It is the check standard beside the reading under test.

6. Model metrology

Every model whose output enters the record is a measuring instrument, and Standard treats it as one. This applies first to the transcription and translation chain, which is already on the record, and to any later inference class without change.

Calibration record. Every model version and every feature-extraction version has one: identity, version, configuration hash, validity domain (languages, acoustic conditions, populations), the reference set it was calibrated against, the date, the result, the uncertainty budget, and the promotion decision with its comparison report. Every inference references the calibration record in force when it was produced; an inference without one is rejected.

Reference cases. A curated, versioned corpus of certified reference cases — sourced only where a basis exists: consented staff recordings, licensed public corpora, synthetic generation for acoustic strata; never callers — stratified across languages, accents, sex, age, communication styles, speech disabilities, neurodivergent patterns, pitch, microphones, radio codecs, noise, sirens, translation-layer audio, synthetic voice. Each case carries its expected outputs, human-panel labels with inter-rater agreement as the label's uncertainty, and the panel's composition. Re-run on a schedule and on every deployment, charted per stratum as a control chart. The extraction layer has its own reference set of test signals passed through the agency's actual audio chain, so a headset or codec change is caught before it looks like a change in anyone.

Uncertainty and the decision rule. An inference carries an expanded uncertainty from a stated budget — extraction repeatability, model reproducibility, input quality, validation coverage for the stratum, reference-label uncertainty — at a stated coverage factor. Any threshold is applied as a decision rule with a guard band: a deviation flags only when it exceeds the threshold by more than the expanded uncertainty; inside the guard band it is indeterminate and recorded as such. A noisier measurement produces fewer flags, not more.

Longitudinal record of the instrument. A new version never overwrites the old version's outputs. On promotion the new version is run over a held historical validation set and both outputs are stored, so "what changed in the population" and "what changed in our measurement of it" are two queries.

Shadow deployment and promotion. A candidate runs silently on live inputs beside the version in service for a ratified period, its outputs stored as shadow and never delivered. The promotion gate compares reference agreement per stratum, output rates on the same population, per-group rate change, and interval distributions against ratified tolerances. A candidate outside tolerance does not promote until a named person records why. Promotion is a recorded human decision with the report as its basis; the promoted version's first calibration record is its as-left state.

Out of tolerance. When a reference check fails or a challenge is upheld, the instrument is out of tolerance for the affected strata from the last passing check. Every output produced under it in that interval is marked *suspect*; every downstream signal is *withdrawn* by a withdrawal event, never deleted; every person affected is notified with the reason. This is the impact assessment a calibration lab runs when a standard is found out of tolerance.

Validity status and the kill switch. Each (language, acoustic context, population stratum) has a validity status on the calibration record — *validated* with current evidence, *not validated*, or *suspended*. Live output is produced only for validated strata in control; any other state writes a refusal: **NO VALID INFERENCE**, with the stratum and the status. Not a lower confidence, not a footnote. The agency can suspend all live inference as a recorded act, and reinstatement is a recorded act with the evidence.

Disparity across the chain. Seven links — source data, signal extraction, model inference, threshold, flag, human response, outcome — each a record class Standard holds, each measurable for comparison groups the agency and its bargaining unit define and the classification register permits it to hold. A difference at any link is reported as a *disparity* with its magnitude and the first link at which groups diverge; bias is a human finding after the investigation a disparity opens. Because the human response is a record, an implementation that answers identical signals differently for different groups is a disparity at link six with the earlier links flat, and the report distinguishes it from a model finding because the chain does. Disparity series are kept over time and intersected with organizational event records and calibration records, so *when did this start* is a query.

7. Review and quality

7.1 The timeline review

Every event in ingest order and in reconciled chronology with intervals, per thread and merged on request, in as-known mode by default for a chosen principal and decision, with the recorder's audio and screen one click away by call identifier. Recorder vendors reconstruct audio, screen, CAD, and radio today; what the edition adds is the AI artifact beside the human action, the joined cross-vendor record independent of any one

vendor, the knowledge state with its capture grades, and the review findings appended to the same record as the evidence.

7.2 The variance lifecycle

A **variance** is a signal, never a conclusion: expected (from protocol outputs, policy, or configuration, with the version), observed, delta. Detection is deterministic and reproducible, which is what makes sampling defensible. The lifecycle is *detected* → *contextualized* → *explained* → *reviewed* → *finding*, each step its own event with its own actor and timestamp, each visibly after the decision it concerns. The **context packet** is assembled automatically and frozen: the actor's knowledge state in as-known mode with capture grades; concurrent threads on the occurrence; concurrent workload at the position; field inputs that had arrived; policy and protocol versions in force; integration and platform health at the time; the recorder reference. The telecommunicator's explanation is their own account, marked post-hoc. The finding is one of a set the agency defines and can be superseded, never erased. Variances are never a performance statistic; My Metrics shows findings by category and the explanations given. Protocol logic stays with its owner: variance is measured against the protocol system's outputs and the agency's own policy, and the protocol vendor's compliance scores, where used, arrive as external assertions.

7.3 The reconstruction report

In the Director's order, every line graded: what we received (the receipt ledger, with gaps stated as gaps — "no supplemental-data events 02:14–02:31; undetermined whether none were sent or delivery failed"); what each system said (threads unmerged, discrepancies listed); what each person had (knowledge state per position with capture grades); what protocol recommended, policy required, the system recommended, the human decided, and the field did (the decision chain with authority, *absent* where a layer was not captured); where they diverged (variances at their lifecycle status, nothing promoted past it); what we can prove (the grade summary, intervals on time claims, indeterminate order named). Rendered through the deterministic vocabulary projection and hash-verified against its events, so it reproduces in five years and can be compared. CAD remains the original for CAD data; the report cites it rather than impersonating it. Hash verification supports self-authentication of the platform's own output; it shortens the custodian's affidavit rather than eliminating it.

8. Operating the agency

The shift. The schedule is built against a required-staffing curve by hour and day. Personnel clock into a console position, not just the building; timecard and position assignment are one record. Position assignment is gated on certifications — EMD, EFD, EPD, CPR, state licensure, CJIS training — with expiry, so the agency can prove who was certified for what on any date. Overtime, trades, and leave run through the workflow engine with the agency's approval chains.

Live coverage. Three dimensions: unit and apparatus coverage as a projection of unit-status events; console and staffing coverage against positions required, now and for the next hour; data coverage — of today's incidents, how many received device-derived location, a supplemental profile, or video — with integration liveness turned into a business signal by integrity monitoring, so a feed that silently stops is noticed before a week goes without it.

Incident flow. The pipeline board over CAD threads — created, assigned, en route, on scene, transporting, closed — with age-in-stage; the persistent dock keeps the active-incident table visible across screens.

Post-incident process. Closure of a defined incident type triggers the agency's follow-up checklist with owners and deadlines; a hand-off to a neighboring agency opens a task that closes only when that agency's disposition arrives or escalates; the workflow's own history is on the record.

Equipment. Apparatus and life-safety devices carry inspection and calibration schedules; a device coming due is pulled from the assignable pool and a work order opens; used on a call, its compliance state at that moment is on the incident. The equipment belongs to the fire or EMS agency's tenant, the unit status to the dispatching ECC's, joined by the same relationship model that links mutual-aid partners.

Records and documents. Incident timeline reports, court packets, council reports, and public-records responses, each hash-verified. Records requests are cases with statutory deadlines, redaction profiles, and delivery events; the structured record and its documents are in scope, and audio and video redaction stay with the recorder, referenced by the case.

Relationships and communication. Book of business becomes served agencies, mutual-aid partners, the fire and EMS agencies whose equipment is linked, and the vendors whose data is consumed. Chat attaches to the incident, the review, the apparatus.

9. Human state — conditional

The Director's questions eight and nine describe inference of regulation state for callers and known users, longitudinal change from self, and hard governance. The design exists (the round-three response) and is summarized here because the edition must be able to carry it. It is **not scheduled**, and it is not built before four things exist in her words: a legal review (some jurisdictions prohibit workplace emotion inference outright and others regulate vocal feature extraction from anyone, and callers cannot consent); the off-limits ruling on whether her center wants any of it; consent and bargaining-unit terms for known users; and a ratified purpose register. Her ruling it out entirely is a legitimate outcome.

If built, it inherits everything above and adds five record classes — observed signal (numbers with units and intervals; audio never held), inference (a closed dimensional vocabulary with an interval and a calibration record), refusal (reason codes; **INSUFFICIENT EVIDENCE — NO INFERENCE** is a row), context, and interpretation (a human, post-hoc, by name) — with speaker attribution as its own inference, trajectories rendered as bands, intersections stated as association with lag and never cause. The firewall between state signal and performance finding is four layers deep: a purpose register with prohibited purposes present only so refusals can be audited; a separate schema with a fence asserting no join to personnel; writer rules that reject any finding citing an inference and any metric derived from state; and role classes in which QA, HR, and fitness-for-duty carry nothing. Longitudinal analysis is opt-in and revocable, compares only to the person's own instrument-versioned baseline, reports a factual exposure ledger beside every signal, runs the system and team checks before any individual signal and freezes baselines while a team-level deviation is open, emits only a fixed-template signal with nowhere to write a diagnosis, triggers only a check-in invitation, and delivers to the employee first. A wellness signal never outranks the human being.

10. Scale and availability

A large county center handles on the order of a million calls a year; the largest metros several million to roughly ten million. With every state change captured, a large tenant produces on the order of one hundred

million to one billion incident events a year. Referencing rather than ingesting media and clinical content is what keeps storage in the terabytes at the top of that range rather than beyond it.

Incident events live in their own schema on the dedicated tier, partitioned by tenant, month, and record class. Append-only storage has no update churn; vacuum is freeze maintenance; retention is a partition drop with extraction and key destruction; recent partitions are hot and small, history cold and never written. Search over events is on structured fields served by tenant-leading indexes; free-text search over transcripts is not offered in this edition. Write throughput is not the constraint: a stadium surge multiplying call volume more than tenfold within an hour produces hundreds of events per second, within a single PostgreSQL primary. The constraints are vendor rate limits, webhook delivery and retry, and idempotent replay.

The edition is not in the call path and is not priced as if it were. Targets: ingest zero-loss, the queue as the availability boundary with seven days' retention; 99.9 % monthly for interactive use; the dedicated tier on synchronous multi-AZ replication in-region (RPO 0) with an asynchronous cross-region replica for disaster recovery (RPO ≤ 1 s, RTO ≤ 15 min); the shared tier multi-AZ. Synchronous cross-region replication is deliberately not the design.

11. Compliance program

- **CJIS.** Background checks, access controls, FIPS-validated encryption, and audit meeting a state-administered review; hosting in a CJIS-eligible region; the classification register and the class-scoped audit as what the review certifies; a state-scoped shared tier where a CSA refuses co-tenancy, or the dedicated tier. Starts before the first ECC pilot.
- **HIPAA.** Broker medical profiles and EMS data under the PHI class, with a business associate agreement path before the broker adapter stores medical payloads.
- **Records law.** Retention, 911-recording exemptions, and redaction by state and record type; the policy model handles the variance and the content is researched per jurisdiction before each state's first pilot.
- **Data-provider terms.** Verified per adapter before it ships; they decide ingest versus reference per element.
- **Consent and labor.** For any analysis of a known user's voice or interaction pattern: a legal review by jurisdiction, consent scopes as records, bargaining-unit terms, and the purpose register, all before build.
- **Model governance.** Calibration records, reference sets, shadow promotion, and the disparity chain as the standing evidence that every model in the record is fit for its purpose.

12. Technology

Kotlin on JDK 25 LTS with Spring Boot 4.1; sealed types for the handled event model with an explicit unrecognized class. Raw JDBC through a hand-rolled safe-SQL kit, no ORM; every query explicit and inspectable, which is what an auditor asks for and what keeps plans predictable under row-level security. PostgreSQL 18, managed; RLS enabled and forced on every tenant-scoped table with class predicates for classified columns; leakproof-only operators on those columns; tenant-leading indexes; tenant context through async execution. Two tenancy tiers behind a placement lookup. Content-addressed encrypted blob storage under a per-tenant, per-class, per-period key hierarchy. Deterministic vocabulary projection with hash verification for every rendered document. Typed domain events, durable queue with idempotent replay,

in-order writer, tenant propagation throughout. React front end with an instrument aesthetic, functional-only color, tabular numerals.

13. The method

Praxis Discovery comes first and is ten working days. The Director is the domain expert: two to four interview sessions across the eight worksheets — personas, records, workflows, standards, numbers, vocabulary, off limits, incumbents — written in her words, with "show me" until a real form or report is on the table; three days of research with every claim tagged observed, claimed, or inferred; two days for the System Picture; one for the foundation and the stack; one for the wedge; one for the grill. The Picture's spine section is Section 4 of this document corrected into her words. Its foundation answers are Standard's defaults with two conditions triggered — a mandated compliance framework, and residency where a state requires it — and its stack path is a pack on Standard unless the grill finds a spine mechanic the substrate cannot carry.

What she supplies, consolidated from three rounds, in worksheet order: the names for the occurrence, the call, the CAD record, the protocol session, and the review; whether the reconstructed timeline is the spine; the decision points and who decides; the response plan and which layers her CAD exposes; what each seat's screen holds and what is shown versus available; the authority register; the location precedence as practiced; the QA form, one completed example, the named exceptions, and the finding categories; the AI services in use and the interpreter practice; whether a Standard panel at the position crosses the line; the recorder and whether it exports; the mutual-aid and CAD-to-CAD arrangements and what her recorder captures; the relations her CAD, RMS, and fire RMS produce; the retention schedule by class and who may order sealing and expungement; which elements are criminal-justice information under her CSA's reading, the roles and what each may see, and what the CSA has said about cloud and co-tenancy; the full off-limits list; and, only if the law and her judgment permit, the terms under which any known user's voice is analyzed and the comparison groups her bargaining unit considers legitimate.

Her confirmation is an input recorded against a brief; ratification stays with the owner. Then the wedge is built in thirty days with its acceptance audit written first — isolation, the spine's walls, the fences watched biting, the documents rendering, the persona's three questions answered on a real screen, the export working, the design partner's number moving — and the record that answers her questions is built behind it, on the primitives the wedge already proved.

14. Why this works

Standard was not a workflow tool that grew an audit log, and it was not an audit log that grew a scheduler. It was built as an immutable record spine with the entire operation of an industry running on top of it, by a method that tags every claim by how it is known and refuses to build from a guess, in a domain where a record must be reproducible and defensible for years and where the instruments that produce it are calibrated against references and recalled when they drift.

Public safety asked for exactly that, in three rounds of questions, and did not ask for anything else: a record that preserves what was knowable, refuses what it should not know, never turns a guess into a fact, and calibrates the thing doing the measuring. The spine and the tenancy carry over unchanged. The operations layer carries over with extensions. The record model, the metrology, and the review are new.